

1 Introduction

The Pebble Bed Modular Reactor (PBMR) power plant is a helium-cooled, graphite-moderated High-Temperature Reactor (HTR). The reactor has a vertical steel pressure vessel which contains and supports a metallic core barrel, which in turn supports the cylindrical pebble fuel core. This cylindrical fuel core is surrounded on the sides by an outer graphite reflector and at the top and bottom by graphite structures that provide similar upper and lower neutron-reflection functions. As a fuel source, the PBMR uses particles of enriched uranium dioxide coated with silicon carbide and pyrolytic carbon. The particles are encased in graphite to form a fuel sphere or “pebble” about the size of a billiard ball. The core of the reactor contains approximately 360 000 of these fuel spheres.

Helium, which is used as the coolant, transfers the energy absorbed in the core to a secondary loop through a special heat exchanger into a steam generator. The helium in the primary circuit is circulated by a blower. A major operational concern is a possible air ingress scenario whereby oxygen would enter the reactor core via a rupture or leak in a pipe. During such an incident the graphite would oxidise rapidly, with potentially disastrous consequences. Thus it is important to have a clear understanding of the kinetics that govern the oxidation of the graphite under these conditions, in order to predict the extent to which a thermal runaway is possible.

The high-temperature oxidative behaviour of graphite is of concern in many other applications due to its wide and varied industrial use. For example, due to its relatively inert nature, graphite is used as an electrode material and a refractory in the metal smelting industry. In this application its longevity is directly related to its oxidative stability under very harsh environmental conditions. This is equally true of applications where graphite is used as a high-temperature gasket material. Thus an understanding of graphite oxidation on a fundamental level could lead to benefits in a wide variety of industrial applications.

Despite large amounts of research on the oxidation of carbon and graphite materials, much uncertainty still surrounds the fundamental aspects of the oxidation. This is largely related to a lack of clarity regarding the exact kinetic mechanism whereby oxidation proceeds. In addition, very few studies have focused on seeking an understanding of the governing mechanisms across the

entire conversion range. Instead, most work has been focused on the initial region of the oxidation curve as a basis for comparing graphite samples from different origins.

Due to the large number of complex mechanisms that affect the oxidation of graphite, this approach is simply not sufficient. There is a need for a structured approach to decouple the relevant effects systematically. This will allow the comparison of different samples on an equal footing. For example, the effect of micro-structural differences between natural and synthetic graphite materials could then be compared, whereas without this framework these differences may be overshadowed by the effect of catalysts present in the natural graphite samples – a fact that would be overlooked if one simply compared the initial reactivities at the same temperature.

It is the purpose of this investigation to evaluate and separate the mechanisms that affect graphite oxidation. The relevant mechanisms can then be identified and an understanding of their operation across the entire conversion range can be developed. This understanding will allow the formulation of a coherent model which can be used as the basis for comparing the relative reactivities of samples from different origins.

A large range of different graphite samples from both synthetic and natural origins were evaluated in this study. This was done in order to explore the wide variety of phenomena encountered during oxidation. Only oxidation under atmospheric pressure in pure oxygen was considered and comparatively low temperatures (<900 °C) were chosen to avoid mass transfer limitations. To further avoid any possible mass transfer limitations, only powdered samples were used in this study. A variety of characterisation techniques were employed and an extensive microstructural investigation was done using ultra-high resolution scanning electron microscopy. All kinetic investigations were done using thermogravimetric analysis (TGA). Modelling and simulation were done in Microsoft Excel® and on the Matlab® platform.

2 Theoretical background

2.1 Atomic structure of carbon

Since the days of Rutherford and Bohr it has been a well-established fact that the atom consists of a positively charged nucleus surrounded by negatively charged electrons. The nucleus is composed of one or more protons and neutrons. Together they account for practically the entire mass of the atom. Each proton has a positive electrical charge and the neutrons are electrically neutral. The electrons move around the nucleus and carry a negative electric charge, equal and opposite to that of a proton. For each proton in the nucleus one electron orbits it, giving an atom that is electrically neutral in its ground state.

As determined by Schrödinger, the orbit of an electron around the nucleus is governed by the specific rules of standing waves. That is to say, the amplitude to find the electron at different places in space can be represented by a function, ψ (referred to as the wave function), of position in space and time. As a non-relativistic approximation and assuming constant angular momentum, the movement of charged particles in an electrostatic field obey the following expression [1]:

$$i\hbar \frac{\partial \psi}{\partial t} = \hat{\mathcal{H}} \psi \quad (2.1)$$

with

$$\hat{\mathcal{H}} = -\frac{\hbar^2}{2m} \nabla^2 + V(r) \quad (2.2)$$

Here $\hat{\mathcal{H}}$ is the so-called Hamiltonian operator, which generates the value for the total energy of the system and depends on the specific situation under consideration. In this case, m is the mass of the electron and $V(r)$ is the potential energy of the electron in the electrostatic field of the proton (which is assumed to be stationary).

The solution of this expression leads to the currently accepted model of the atom that governs nearly all chemistry. This model states that, in any given atom, the electrons are found in a series of discrete energy levels, which are known as orbitals. These are distributed around the atom in well-defined regions

of space where there exists a high *probability* of finding the electron. Some illustrative representations of these orbitals are shown in Figure 2-1.

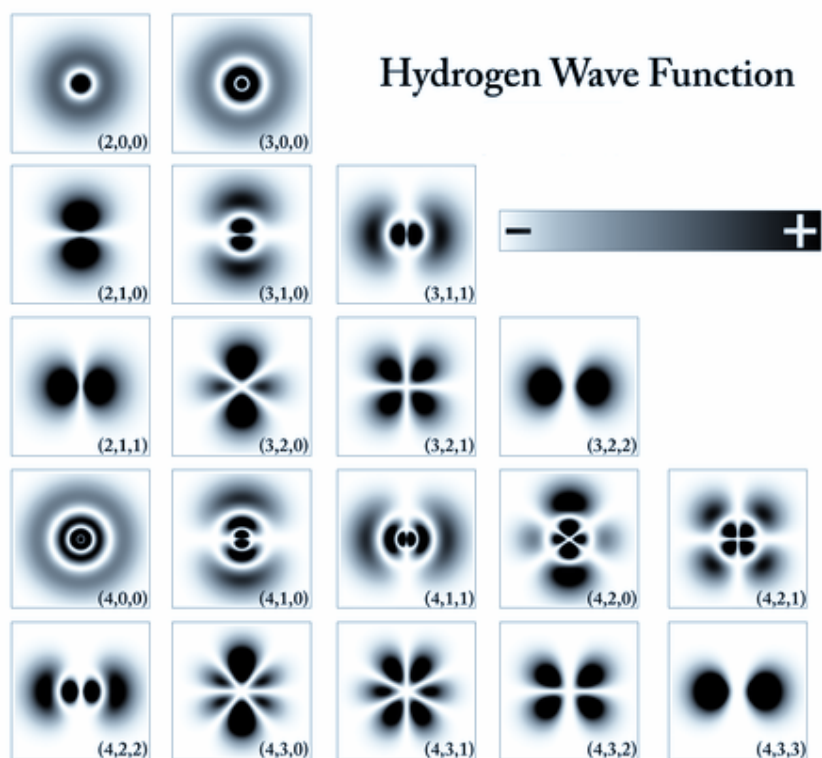


Figure 2-1: Solutions to the hydrogen wave function showing orbital configurations

As dictated by the Pauli exclusion principle, no more than two electrons, each with opposite spin, can occupy any given orbital. The element carbon (symbol C) has an atomic number (Z) of six. Thus the neutral atom has six protons in the nucleus and, correspondingly, six electrons. For the most abundant, naturally occurring, stable, carbon isotope, the nucleus contains six neutrons. This gives carbon a mass number (A) of 12.

The six electrons of carbon are arranged in a $1s^2 2s^2 2p^2$ configuration, as dictated by Eq. (2.1). The notation $1s^2$ refers to the quantum numbers necessary to define an orbital. The number one (principle quantum number) refers to the first shell of electrons (K shell) and the letter "s" (angular momentum quantum number) refers to the type of subshell or orbital. The shape of an s-orbital is shown in Figure 2-2. The superscript two refers to the number of electrons in that

subshell. The K shell only has one orbital (the s-orbital), which can never have more than two electrons. These electrons will have opposite spin, are closely bound to the nucleus and have the lowest possible energy. The filled K shell is thus completely stable and hence electrons from this shell do not participate in chemical bonding.




Figure 2-2: Schematic representation of an s-orbital

The next two terms ($2s^2 2p^2$) refer to the four electrons in the second shell or L shell. This shell has additional allowable energy levels in the orbital configuration known as the “p” orbitals. The shape of a p-orbital is shown in Figure 2-3. There are three possible p-orbital configurations in three dimensions, namely x, y and z.




Figure 2-3: Schematic representation of the p-orbitals

Hence this shell is filled when the atom has eight electrons, as is the case for the noble gas neon. In such a configuration the L shell is completely filled and the atom is stable and chemically unreactive. Carbon, on the other hand, has only four electrons in its outer shell and seeks to make chemical bonds with itself and other elements to fill this shell. Due to differences in the energy levels of the s and p-orbitals, the s-orbital is filled first, followed by one electron in two of the three p-orbitals. This electronic structure can be represented schematically as shown in Figure 2-4 and Table 2-1.

Table 2-1: Electronic structure of carbon in the ground state

K-shell electrons	L-shell electrons			
1s	2s	2p _x	2p _y	2p _z
↑↓	↑↓	↑	↑	




Figure 2-4: Schematic electronic structure of carbon

In any given atom the electrons located in the outer orbital are the only ones available for bonding with other atoms and are known as the valence electrons. Hence for carbon these are the two 2p electrons and carbon is divalent. Divalent carbon does exist, such as the highly reactive, transient-organic intermediates: carbenes (e.g. methylene). In 1991 a stable, crystalline dicarbene was isolated by Arduengo *et al.* [2]. However, the most abundant, stable, carbon allotropes are not divalent but tetravalent. The reason for this is orbital hybridisation. It is this phenomenon that produces the large variety of carbon allotropes that exist. More detailed accounts of the concepts in this section and the following section can be found in many standard chemistry textbooks, e.g. *Atkins' Physical Chemistry* [3].

2.2 Carbon allotropes

The word carbon is derived from the Latin *carbo*, which to the ancient Romans meant charcoal (or ember). In the modern world carbon is much more than simply charcoal. This is because, unlike most elements, carbon has an almost limitless variety of allotropes (or polymorphs). These materials are

composed entirely of carbon but have vastly different physical structures and properties. A few of these allotropes are shown in Figure 2-5, namely a) diamond, b) graphite, c) Lonsdaleite, d) C₆₀ (Buckminsterfullerene or buckyball), e) C₅₄₀, f) C₇₀, g) amorphous carbon, and h) single-walled carbon nanotube.




Figure 2-5: Some carbon allotropes

These allotropes are the result of the wide variety of bond structures that carbon atoms can form with each other. A chemical bond is formed when an electron comes sufficiently close to two positively charged nuclei, such that it is attracted by both simultaneously. This form of chemical bonding is known as covalent bonding and is the strongest type of chemical bonding. In order to achieve the lowest possible energy configuration, the s and p-orbitals can combine to form hybrid orbitals. These orbitals, in conjunction with the chemical bonds between atoms, i.e. electron sharing, yield the lowest energy configuration for that specific molecule's atomic composition. The simplest example of this phenomenon for carbon is its chemical bonding with hydrogen. The carbon atom can form three types of hybrid orbital, namely four sp^3 -orbitals, three sp^2 -orbitals or two sp -orbitals. These orbitals form direct covalent bonds with their

neighbours, leading to the three types of carbon-hydrogen molecule shown in Figure 2-6.




Figure 2-6: (1) Ethane C_2H_6 : sp^3 -hybridisation; (2) Ethene C_2H_4 : sp^2 -hybridisation; (3) Ethyne C_2H_2 : sp^1 -hybridisation

The first type of hybridisation occurs when all four orbitals of the L shell are combined to form four sp^3 -orbitals. A visual representation of the sp^3 -hybrid orbital is shown in Figure 2-7. The electrons in the L shell are evenly distributed across these four orbitals as indicated in Table 2-2.

Table 2-2: sp^3 hybridised electronic structure

K-shell electrons	L-shell electrons				
	1s	2s	2p _x	2p _y	2p _z
↑↓	↑	↑	↑	↑	↑

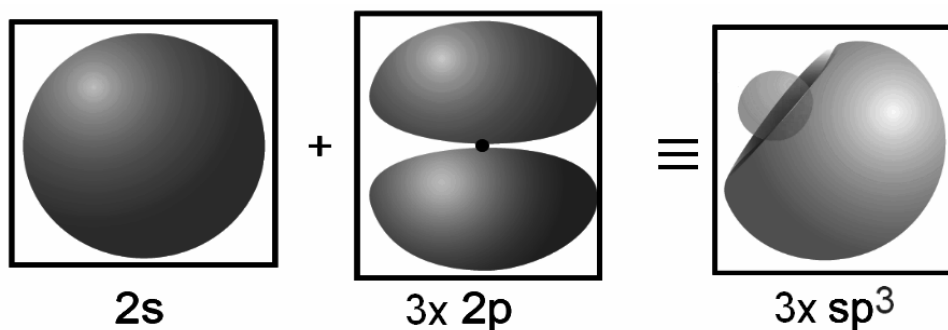


Figure 2-7: sp^3 -hybrid orbital

To maintain the lowest energy configuration, these four orbitals are arranged in a tetrahedral formation with the maximum equal angle (109.5°) between each, as shown in Figure 2-8.




Figure 2-8: Tetrahedral formation of sp³-hybrid orbitals
(negative lobes omitted for clarity)

When such a hybrid orbital overlaps with an orbital from another atom, the nuclei share a pair of electrons and a covalent bond is formed. Due to the small size of the carbon nucleus and the fact that this hybrid structure allows four of carbon's six electrons to form bonds, the resulting covalent bonds are quite strong. By convention these strongly directional (or stereospecific) bonds are called sigma (σ) bonds. This hybridisation is found in all aliphatic carbon compounds. Due to the heavily lopsided nature of these orbitals, a substantial overlap is possible when the orbitals combine with other sp³-orbitals from other carbon atoms. This further enhances the strength of the bonds and allows a three-dimensional, entirely covalent, tetrahedral lattice structure. This arrangement is the basic structural unit of the diamond crystal, as shown in Figure 2-9.




Figure 2-9: Tetrahedral structure of diamond

The second type of hybridisation that is possible for the carbon atom is sp^2 -hybridisation. This arrangement forms the basis of all graphitic structures and aromatic compounds. In this case the s-orbital combines with only two of the three p-orbitals to form three sp^2 orbitals. The remaining p-orbital remains in its unhybridised ground state. Three of the four electrons occupy the three hybrid orbitals, while the final electron exists in a “free” or delocalised state within the remaining p-orbital. A visual representation of the sp^2 -hybrid orbital is shown in Figure 2-10. The electrons in the L-shell are evenly distributed across these four orbitals as indicated in Table 2-3.

Table 2-3: sp^2 hybridised electronic structure

K-shell electrons	L-shell electrons			
1s	2s	2p _x	2p _y	2p _z
↑↓	↑	↑	↑	Free delocalised electron




Figure 2-10: sp^2 -hybrid orbital

To maintain the lowest energy configuration, these orbitals are arranged in a planar, trigonal fashion with the maximum equal angle (120°) between them, as shown in Figure 2-11.




Figure 2-11: Planar trigonal formation of sp^2 -hybrid orbitals

The fourth, unhybridised orbital lies perpendicular to the plane of the three sp^2 -orbitals. The electron in this orbital is capable of forming a significantly weaker bond with adjacent atoms, called a pi (π) bond. The orbitals are similar to the sp^3 -orbitals but in this case even more pronounced overlap is possible with other sp^2 hybridised carbon atoms. This leads to a shorter bond length and higher bond energy in the orbital plane, which implies that in-plane graphite has a higher strength than diamond.

When bonded in plane to other sp^2 hybridised carbon atoms, the 120° bond angle leads to the formation of a planar hexagonal structure. Furthermore, when several of these planes are arranged on top of each other, a layered structure results. This arrangement forms the basic structural unit of the graphite crystal, as shown in Figure 2-12.




Figure 2-12: Hexagonal planar structure of graphite

In the graphitic structure, the delocalised electron has a weak van der Waals interaction with the adjacent layer [4]. This leads to the extremely low in-plane shear resistance of graphite. It is presumed that the delocalised electrons form a resonance structure (or conjugated system) similar to that of benzene, shown in Figure 2-13. This is especially true for a single layer of graphite, called graphene. The delocalised electrons can easily move from one side of the plane to the other, but cannot readily move from one layer to the next. This leads to the anisotropic nature of the physical properties of graphite.




Figure 2-13: Resonance structure of benzene

The final hybridisation that is possible for carbon is the formation of two sp^1 -orbitals. In this case the L-shell s-orbital combines with one p-orbital to form two sp^1 -orbitals which are orientated at 180° to each other. The remaining two p-orbitals remain unhybridised, perpendicular to the direction of the sp^1 -orbitals and with a 90° angle between each other. The structure is shown in Figure 2-14.




Figure 2-14: Schematic representation of sp^1 -hybridisation configuration

Electrons are distributed evenly across the hybridised and unhybridised orbitals. Electrons in the sp^1 -orbitals form strong covalent σ bonds with other atoms, while the two delocalised electrons in the p-orbitals are available to form π bonds in a similar manner to the sp^2 -hybridisation case. These covalent bonds represent the strongest carbon-carbon bond with the shortest bond length, e.g. acetylene. However, in chain form, these carbynes are usually unstable [5].

These three hybridised configurations (or combinations thereof) account for the molecular structure of all the stable carbon allotropes. This includes all forms of char, coal and anthracite, as well as the so-called “nanostructured” carbons such as fullerenes, nanotubes and graphene sheets. Fullerenes and nanotubes may be thought of as three-dimensional derivatives of single graphene

sheets that have been folded into specific conformations, as can be seen in Figure 2-15. These structures have many unique properties and interesting applications, but they are considered to be outside the scope of this investigation. However, two of the carbon allotropes are of particular importance and will be discussed in more detail in the following two sections.




Figure 2-15: Configurations of graphene, fullerenes, nanotubes and graphite

2.3 Graphite

The word graphite is derived from the Greek word *graphein*, which means to write, due to its extensive use in pencils. This use illustrates once of the many interesting and unique properties of graphite, namely its low in-plane shear resistance. These properties are directly related to the atomic structure of graphite as mentioned in the previous section. Ideally, graphite can be considered to be a planar arrangement of stacked, infinite sheets. Each sheet consists of a network of tightly bound carbon atoms in a hexagonal lattice, as represented schematically in Figure 2-16.




Figure 2-16: Graphite lattice

However, in reality this is an exceedingly simplified picture. Categorically defining the structure of graphite as a material is very difficult. This is due largely to the wide variety of different origins of materials that are broadly classified into the general category of graphite. All these materials have correspondingly different micro- and macrostructures, and potentially even different atomic structures. The latter is due to some uncertainty regarding the possible presence of small amounts of sp^3 -hybridised atoms within certain types of graphite.

This section focuses on exploring the properties of the idealised graphite structure, as well as a few of its readily obtainable natural and synthetic forms. Finally, some background is provided on the analysis techniques that can be used to characterise graphite samples with different origins and to identify their key distinguishing traits relevant to oxidation.

2.3.1 Graphite crystal structure

Crystallography covers the arrangement of atoms in very a wide variety of solid materials. It is therefore necessary to define a standard structure within which different crystal structures can be characterised and differentiated. A brief classification of crystal structures will be given here. For a more detailed analysis the reader is directed to Ashcroft and Mermin [6].

A crystal is made up of a periodic arrangement of one or more atoms which repeat at every lattice point. Consequently, the crystal looks the same when viewed from any of the lattice points. A crystallographic point group is a set

of symmetry operations, such as rotations or reflections, that leave a point fixed while moving each atom of the crystal to a position of an atom of the same kind. That is to say any infinite crystal would look the same before and after any of the operations in its point group. There are infinitely many three-dimensional point groups; in crystallography, however, they are restricted to 32 to be compatible with the discrete translation symmetries of the known crystal lattices. When these point groups are combined with the 14 Bravais lattices, the complete set of space groups for crystallography results.

A lattice system is a class of lattices with the same point group, of which there are seven. The lattice system of a crystal or space group is determined by its lattice but not always by its point group. The 14 Bravais lattices are obtained by uniquely combining the seven lattice systems with one of the four possible lattice centerings. The lattice centerings are primitive centering, body centred, face centred and base centred. This is one unique way of defining crystal structure.

An alternative characterisation is via the crystal family and system. A crystal system is a class of point groups. Two point groups are placed in the same crystal system if the sets of possible lattice systems of their space groups are the same. For many point groups there is only one possible lattice system, and in these cases the crystal system corresponds to a lattice system and is given the same name. However, for the five point groups in the trigonal crystal class there are two possible lattice systems: rhombohedral or hexagonal.

A crystal family also consists of point groups and is formed by combining crystal systems whenever two crystal systems have space groups with the same lattice. In three dimensions a crystal family is almost the same as a crystal system (or lattice system), except that the hexagonal and trigonal crystal systems are combined into one hexagonal family. These classifications are most easily visualised by the use of a table, as shown in Table 2-4.

Table 2-4: Crystal classification

Crystal family	Crystal system	Required symmetries of point group	Point groups	Space groups	Bravais lattices	Lattice system
Triclinic		None	2	2	1	Triclinic
Monoclinic		1 two-fold axis of rotation or 1 mirror plane	3	13	2	Monoclinic
Orthorhombic		3 two-fold axes of rotation or 1 two-fold axis of rotation and two mirror planes	3	59	4	Orthorhombic
Tetragonal		1 four-fold axis of rotation	7	68	2	Tetragonal
Hexagonal	Trigonal	1 three-fold axis of rotation	5	7	1	Rhombohedral
	Hexagonal			18		Hexagonal
Cubic		4 three-fold axes of rotation	5	36	3	Cubic
Total:	7		32	230	14	7

A unit cell may be defined as the smallest, regularly repeating material portion contained in a parallelepiped from which a crystal is formed by parallel displacements in three dimensions. That is to say it is the smallest possible representation of the entire lattice. The generic unit cell is shown in Figure 2-17.




Figure 2-17: Basic unit cell

The definition of the unit cell for each of the seven crystal systems is given in Table 2-5. Notice that some trigonal lattices can be expressed on the basis of either a hexagonal or a rhombohedral lattice.

Table 2-5: Crystal unit cell classification

Crystal system		Cell parameters
Triclinic		$a \neq b \neq c;$ $\alpha' \neq \beta' \neq \gamma'$
Monoclinic		$a \neq b \neq c;$ $\alpha' = \gamma' = 90^\circ, \beta' \geq 90^\circ$
Orthorhombic		$a \neq b \neq c;$ $\alpha' = \beta' = \gamma' = 90^\circ$
Tetragonal		$a = b \neq c;$ $\alpha' = \beta' = \gamma' = 90^\circ$
Trigonal	Hexagonal	$a = b \neq c;$ $\alpha' = \beta' = 90^\circ, \gamma' = 120^\circ$
	Rhombohedral	$a = b = c;$ $\alpha' = \beta' = \gamma' \neq 90^\circ$
Hexagonal		$a = b \neq c;$ $\alpha' = \beta' = 90^\circ, \gamma' = 120^\circ$
Cubic		$a = b = c;$ $\alpha' = \beta' = \gamma' = 90^\circ$

Miller indices are a notational system developed for designating planes and directions in crystallography using Bravais lattices. A particular family of lattice planes is determined by three integers, h , k and l (the Miller indices). Each index denotes a plane orthogonal to a direction in the basis of the reciprocal of the lattice vectors, as shown in Figure 2-18, where a , b and c are the dimensions of the unit cell as defined before. Hence the lattice planes designated by the reciprocal of the Miller indices are multiples of the unit cell.




Figure 2-18: Miller indices and corresponding planes

The graphite crystal falls into the hexagonal crystal family. The ball and stick model of the most common graphite crystal system and lattice is shown in more detail in Figure 2-19. The figure includes the standard definition of the unit cell, with the relevant dimensions given in Table 2-6.

Table 2-6: Hexagonal graphite unit cell data [7]

Crystal system	Hexagonal
Space group	d_{6h}^4 - $P6_3/mmc$
a (Å)	2.464
b (Å)	2.464
c (Å)	6.711
α' (°)	90
β' (°)	90
γ' (°)	120
Calc. density (g/cm ³)	2.26
Unit cell volume (10 ⁶ pm ⁻³)	35.29




Figure 2-19: Hexagonal graphite crystal structure

The in-plane sp^2 -hybridised atoms are bound to three adjacent neighbours to form the characteristic hexagonal structure, with an internal angle of 120° . This plane is commonly referred to as the basal plane of graphite. These σ bonds are short, 0.141 nm, and have a high strength, as indicated by their high bond energy, 524 kJ/mol. These atoms can be grouped into two categories, i.e. those that have neighbours directly above and below (alpha) in adjacent planes and those that do not (beta).

The unhybridised fourth electron is delocalised and interacts with the corresponding delocalised electrons from the adjacent plane. It seems unlikely that each electron has a direct interaction with a neighbour in the adjacent plane since the beta atoms do not align with any direct neighbours. Instead, as mentioned earlier, the delocalised electrons probably form a resonance-based electron cloud, similar to that of benzene, but spread across the entire plane. This cloud interacts with the adjacent layer to form a weak van der Waals-type bond due to structural polarisation. The bond energy is estimated as 7 kJ/mol.

The interlayer spacing of graphite is fairly large, 0.3354 nm, which is roughly twice the van der Waals radius of the carbon atom. Carbon is the only element that forms this layered hexagonal structure. The large distance between the adjacent layers allows the intercalation of a variety of ions and molecules between the graphite layers. This has led to the development of a wide variety of GICs (graphite intercalated compounds) which have a range of interesting applications, including lithium-ion batteries and fire retardants.

However, two potential stacking arrangements exist for the graphite crystal structure. The first is the hexagonal arrangement that is shown in Figure 2-19. The stacking order for this form is given by $-ABABAB-$ and thus the layer alignment repeats every second layer. When observed from above the sheets align as shown in Figure 2-20.




Figure 2-20: Hexagonal crystal stacking

Another stacking arrangement, namely the rhombohedral form, is possible for graphite crystals. In this arrangement the third graphitic layer is offset from the first to form an $-ABCABC-$ stacking sequence, as shown from above in Figure 2-21.




Figure 2-21: Rhombohedral crystal stacking

Despite belonging one of the seven space groups associated with the rhombohedral lattice system, the unit cell of rhombohedral graphite is the same hexagonal unit cell used for hexagonal graphite with the exception that the c-direction spans three layers of graphite instead of just two, as can be seen in Figure 2-22. The figure also includes the rhombohedral lattice structure for clarity. The unit cell data for this configuration are shown in Table 2-7.

Table 2-7: Rhombohedral graphite unit cell data [8]

Crystal system	Trigonal
Space group	$D_{3d}^5-R\bar{3}m$
a (Å)	2.566
b (Å)	2.566
c (Å)	10.062
α' (°)	90
β' (°)	90
γ' (°)	120
Calc. density (g/cm ⁻³)	2.28
Unit cell volume (10 ⁶ pm ⁻³)	52.47




Figure 2-22: Rhombohedral crystal structure

This form of graphite is thermodynamically unstable and can be considered as a type of extended stacking fault. During heat treatment above 1 600 K this form reverts to the hexagonal form [8]. Due to the fact that the stacking fault energies in general are quite small, the occurrence of stacking faults in graphite is quite frequent and a very wide variety of faults is possible, as discussed by Amelinckx *et al.* [9]. A related effect is crystal twinning, or twin lamellae, which is highly prevalent in graphite. It is widely accepted that a twin may be derived from the parent graphite crystal by a rotation about the $\langle 10\bar{1}0 \rangle$ or armchair direction [10]. A twin involves rotation of the basal axis by approximately 20° along a carbon-carbon bond and these usually occur in pairs, effectively forming a kink band, as illustrated in Figure 2-23. Twins can be easily produced by mechanical deformation and are readily discernible in any form of microscopy [9]; [10].




Figure 2-23: (a) Side geometry of twin boundary; (b) Twin band consisting of two twin or tilt boundaries AB and CD

Aside from the rhombohedral and other stacking faults, graphite also has a very wide variety of possible crystal lattice imperfections. These include the possibility of vacant sites, where an atom is missing from the basal plane, or disinclinations, which cause adjacent planes to be no longer parallel. Rotations between adjacent sheets are also possible but perfect alignment of two adjacent sheets is not. Some imperfections are probably caused by growth defects and include non-basal screw and edge dislocations, as shown in Figure 2-24.




Figure 2-24: Non-basal screw and edge dislocations

To identify non-basal screw and edge dislocations it is necessary to cleave an oxidised crystal. When cleaved, screw dislocations will result in steps emanating from the oxidation pits; the latter will be nucleated at the screw. Non-basal edge dislocations, on the other hand, can be identified using Frank's theorem regarding pit spacing and the absence of any steps after cleavage [11]. The presence of edge dislocations is usually characterised by the presence of closely spaced edge pits along a certain crystallographic direction. In some cases high densities of non-basal screw dislocations are found in intercrystallite boundaries of pure twist character, as demonstrated by Thomas and Roscoe [12]. This leads to significant reaction anisotropy as these highly defected regions will oxidise far more rapidly than neighbouring regions. Small angle boundaries of both tilt (edge) and twist (screw) components occur most frequently in graphite. These are often associated with the termination points of twin lamellae due to a large angle of misorientation [12].

As mentioned, these lattice imperfections have a direct relevance to the reactivity of different graphite samples. This will be discussed in more depth in subsequent sections. An additional feature of the graphite crystal that is relevant to its oxidative stability is the edge anisotropy. As can be seen from Figure 2-25,

due to the hexagonal nature of the lattice, there are two possible edge arrangements. These configurations are known as the armchair and zig-zag arrangements.




Figure 2-25: Zig-zag and armchair edge sites

Both tunnelling electron microscopy and scanning tunnelling microscopy (TEM and STM) have recently been used to characterise these edges [13];[14]. They have allowed valuable, but preliminary, insights to be gained into potential defect structures at the edges and their effect on oxidation and oxidation catalysts. It has been shown [15] that these sites have differing oxidative stability and that their relative oxidative reactivities depend on temperature. This is due to a difference in their activation energies of approximately 17 kJ/mol.

The twinning phenomenon mentioned earlier acts as a very important crystallographic marker since, as can be seen from Figure 2-23, a twin band always runs parallel to the armchair axis of the graphite basal plane. Thus by determining whether an edge runs parallel or at an angle of 30° to a twin, the edge can be designated as armchair or zig-zag respectively. Furthermore, the internal edge configuration of hexagonal pits, often found in graphite during oxidation, can be determined in this fashion [10], as shown in Figure 2-26. These are designated as parallel and perpendicular pits respectively. In general though, there is no tendency for pits to accumulate in the vicinity of twins [10].

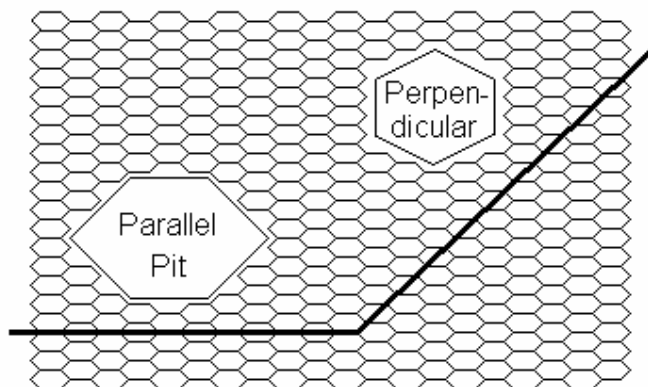


Figure 2-26: Parallel and perpendicular etch pits

It was noted by Thomas [10] that below a certain temperature, around 900 °C, only parallel hexagonal pits are produced in pure, highly crystalline graphite during oxidation. In contrast, above 1 000 °C pits were exclusively orientated in the perpendicular direction. This is a direct consequence of the reaction anisotropy. Below 900 °C the zig-zag edge is more reactive and the armchair conformation dominates, whereas above 1 000 °C the reactivity is reversed and the edges are exclusively zig-zag. It is interesting to note that within the transition region, dodecagonal pits are formed which appear roughly circular in nature and are a combination of armchair and zig-zag conformations.

2.3.2 Natural graphite

Natural graphite is a relatively abundant mineral found in many parts of the world, including, but not limited to, India, China, Brazil, Zimbabwe, Canada and Sri Lanka. It consists of graphitic carbon with varying degrees of crystalline perfection. Hence natural graphite may be further subdivided into micro- and macrocrystalline forms depending on the crystallite size. Microcrystalline graphite, which is also known as amorphous graphite, has very low crystallinity and generally high levels of impurities [16]. This material is formed as a result of the metamorphosis of coal which has been exposed to high pressure.

Two basic forms of macro-crystalline graphite may be distinguished. Firstly there is vein graphite, which occurs mainly in Sri Lanka [17]. It occurs as sizeable crystalline deposits, in the shape of large lumps and chips. These materials show

high levels of very large crystalline regions and very high purity. Hence they have excellent physical properties, showing the closest approximation to the perfect crystal properties for all types of naturally occurring material. Vein graphite is formed by fluid deposition, i.e. the transformation of oil precursors.

The second discernible type of macro-crystalline graphite is so-called “flake graphite”. This material is found mainly in China, Brazil, Zimbabwe and Canada [17]. Flake graphite is found disseminated in metamorphosed silica-rich quartzites, gneisses and marbles. The graphite content in the raw ore can vary from 5 to 40%. Hence these materials contain a significant amount of widely varying impurities. For this reason most naturally occurring graphite is processed further.

After mining, depending on the application, flake and vein graphites are milled. This is usually done with ball mills, although ring rollers and jet mills are also used. The material is then further purified by mechanical separation and flotation to obtain purities in excess of 99.5% carbon. In order to achieve higher purities, >99.9%, chemical and high-temperature treatments are usually required. These include acid washing [18] and alkali roasting [19]. However, these treatments may damage the graphite and leave residual deposits themselves.

A high percentage of the flake graphite production is utilised in the fabrication of refractory crucibles for the metal industry. Other applications include lubricants, brake linings, batteries, electric brushes and pencil lead. Recently an increasing amount of natural graphite is being used to produce GICs. These materials are then exfoliated (allowed to expand at high temperature) and either utilised directly in battery systems or recompressed to form highly conductive graphite sheets (“Grafoil”), which are used in the electronics industry to dissipate heat.

Since natural graphite is not a perfectly crystalline material, its reactivity will depend on the source of the graphite and the extent of graphitisation. This is complicated by the fact that most graphite samples are milled; this destroys some of the macro- and microstructure, leading to the formation of amorphous regions. In addition, naturally mined graphite inevitably contains varying amounts of a wide variety of naturally occurring impurities. The presence of these impurities plays a large role in the reactivity of the natural graphite samples from different origins.

2.3.3 Synthetic graphite

Synthetic carbon materials appear to have their origins in the 1800s when Sir Humphry Davy first employed carbon electrodes to produce an electric arc. Subsequently, in 1896, Castner and Acheson found that some disordered carbons could be converted into graphitic materials. This prompted a new beginning in the carbon industry aimed at producing materials suitable for electrolytic manufacturing, electrothermic production and electric arc furnaces.

Finally, in 1942 when Enrico Fermi and co-workers developed a self-sustaining nuclear reaction using graphite as the moderator, a whole new group of applications for synthetic graphite arose. In recent times there has been a large upsurge of interest in this latter category of synthetic graphite materials. This is due to increasing environmental pressures to move away from heavily polluting coal-driven power generation to greener alternatives.

Today synthetic graphite materials with widely differing properties and physical characteristics are produced. The properties depend on the raw materials used, processing methods, additives and treatments or coatings. Bulk carbon is basically a composite mixture of carbon particles (filler) and a suitable binder. The most common filler is petroleum coke and usually the binder is coal tar pitch. The major processing steps in the manufacture of synthetic graphite are shown in Figure 2-27.

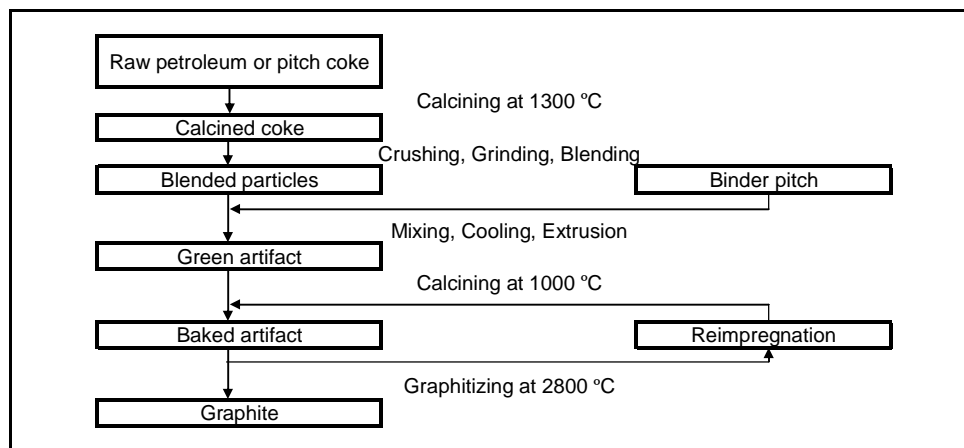


Figure 2-27: Synthetic graphite manufacturing steps

Petroleum coke is obtained largely by the cracking of heavy refining oil to produce gasoline, other light fractions and a heavy product, sometimes known as heavy fuel, decant or slurry oil. This material is then fed to a delayed coking unit to produce green coke. In the delayed coking unit the material is fed at high temperature into a coking drum where it thermally cracks and the coke gradually accumulates in the drum. The formation of coke in the drum involves the transition of the isotropic fluid feedstock, via an intermediate liquid crystal phase (the so-called mesophase), into anisotropic, graphitisable carbon [20].

This process in and of itself is extremely complex and involves several physical and chemical processes. As the temperature is raised, the average molecular weight increases due to polymerisation and condensation. The onset and growth of the liquid crystal droplets can be seen. As these spheres enlarge, they may coalesce depending on the material properties. With continued growth of this anisotropic mesophase, viscosity increases and a resultant fully coalesced structure may result. All the while volatile material is being released, creating macroscopic structural features within the material.

As a result petroleum cokes show remarkable flow patterns within the particles. It has been shown that large plate-like aromatic molecules are aligned parallel to the flow as a result of shear stresses. The extent of this alignment and the consequent microstructure present within the coke will have a dramatic impact on its graphitisation behaviour. This behaviour is governed by the type of feedstock used, as well as the specific conditions during coking. In general, three types of coke are produced, namely sponge or shot, regular and needle. These materials have increasing structural alignment and hence increased graphitisability. Thus they find use in different industries, from reducing agents in iron ore smelting, as fuel or carbon source, to aluminium smelter anodes and electric arc furnace electrodes.

Next the green coke material is carbonised (or pyrolysed). This may be defined as the step in which the organic precursor is transformed into a material that is essentially all carbon. The carbonisation cycle is basically a heating cycle. The precursor is heated slowly in a reducing or inert environment, over a range of temperatures that depend on the objectives and may be as high as 1 500 °C. This is usually achieved by the use of a rotary kiln calciner. However, such

calciners have several disadvantages and in some cases rotary hearth calciners are utilised.

During calcination the aligned structure formed during coking is preserved, although fissures and cracks develop and enlarge as additional volatile hydrocarbons are evolved. Calcined petroleum coke is a turbostratic (disordered) carbon containing crystallites of about 5 nm [21]. During this treatment the coke will undergo considerable shrinkage due to the release of volatiles, which may lead to undesirable results during the subsequent processing steps.

Finally, the coke material is ready for use as a filler. However, to obtain the correct particle size distribution for a specific set of electrode properties, the coke is crushed, milled, sized and blended. The coke can now be mixed with the binder for forming. At this point the correct choice of mixture constituents is critical to ensure the correct degree of anisotropy, density and other properties in the final product.

The binder most commonly used is coal tar pitch, which is the heavy residue derived from the distillation of coal tar, produced as a by-product in coke ovens. The function of the binder is to wet and plasticise the filler so that it can be easily formed into the required shape. Subsequently, the binder must carbonise well during baking. It should give a good carbon yield and form a coke that will bridge the filler particles and form a strong, dense final artefact. This will allow suitable thermal and electrical conductivities to be achieved in the final product.




Figure 2-28: Synthetic graphite forming techniques

Three possible forming techniques are widely used to obtain a variety of product properties, depending on the requirements. These are shown in Figure 2-28. Extrusion is used to form most carbon and graphite products. The extrusion process produces bulk anisotropy in the graphite due to the alignment of the filler particles along the extrusion direction. Moulding can also be used to form

graphite products into more intricate shapes, but it leads to anisotropy within the product material, depending on the type of press used. Finally, the material can also be isostatically pressed; this is the only method for producing a fully isotropic material.

When the formed composite material is baked, the binder will undergo complex transitions, similar to those undergone by the coke. The choice of a suitable binder pitch and the control of these transitions are critical to ensure that the objectives mentioned previously can be met during graphitisation. During baking the diffusion of volatile compounds to the atmosphere is a critical step and must occur slowly to prevent disruption and rupture of the carbon network. As a result, this is a slow process and its duration may vary widely, depending on the type of product being produced; in some cases it may take several weeks.

After it has been baked, the residual material is essentially fully carbonised but with very little graphitic order and is effectively an amorphous material. In some cases this material may be re-impregnated and baked again to reduce porosity and increase density. Depending on the density requirements of the final product, the material may undergo as many as three additional re-impregnation steps. The correct choice of the pitch being used for re-impregnation is critical to ensure complete penetration into the core of the formed shape.

The final step in the production of a graphitic material is the graphitisation step. In general, this is an electrical heat treatment of the baked shape to temperatures around 3 000 °C. An alternative method utilises an induction furnace to heat the material inductively to the required temperature, but this is generally used only for specialist products. The basic elements of the furnace invented by Acheson in 1895 and still used today [20] are a furnace bed consisting of refractory tiles supported by concrete piers.

The product pieces are carefully spaced on a layer of metallurgical coke and the interstices filled with this coke. Most of the heat needed to reach graphitising temperature is generated in this resistive material. This so-called “resistor pack” is covered with a blend of fine metallurgical coke, sand and silicon carbide to provide thermal and electrical insulation. The furnace ends, which are U-shaped concrete heads with several graphite electrodes, are connected by copper bus bar to the secondary of a transformer. A current is passed through the material, allowing heating rates of between 0.7 and 1.0 K/min. Total firing

time is approximately 3 days, followed by a further 10 days of cooling and unloading.

The exact mechanism of graphitisation is still unclear and depends to a large extent on the precursor materials present. The process of graphitisation is essentially the development of an ordered, crystallographic structure from an initially amorphous, but graphitisable, carbon. This is known as turbostratic carbon. Graphitisation occurs as a series of steps which begins when the carbonisation temperature is passed with the release of any non-carbon elements. These are gradually removed up to 2 000 °C, at which point essentially none remain [22]. This transition is shown schematically in Figure 2-29.




Figure 2-29: Transition from turbostratic to graphitic carbon

An extensive study of the structural ordering that occurs during carbonisation and graphitisation for a variety of carbons was conducted by Oberlin [23]. This was done using TEM to characterise the structural development and a variety of analytical techniques, including XRD, Raman and IR spectroscopy, to link the observations to the crystallographic parameters and material properties. Oberlin classified the transition in terms of four distinct stages, which are shown schematically in Figure 2-30. A basic structural unit (BSU) is defined as a planar aromatic structure consisting of fewer than 10 to 20 rings, piled up in a more or less parallel arrangement of two to four layers.

After carbonisation of the mesophase, structures created in the coke and pitch materials already show reasonable local molecular orientation (LMO). The BSUs are small but entirely aromatic, with well-orientated layers. However, they are still fairly randomly distributed within the macrostructure. During this initial stage or stage one, the interlayer spacing within a BSU becomes smaller and more consistent. The next stage, stage two, is characterised by recovery of the columnar order that is present in mesophase single spheres. Small distorted columns which may trap disoriented BSUs are observed and the diameters of the columns do not increase.

Stage three begins after 1 500 °C with the coalescence of the columns taking place, leading to the formation of longer, distorted layers. This stage corresponds to the critical period of edge-to-edge alignment of the columns and the elimination of in-plane defects. Up to the end of this stage the material still remains largely turbostratic. The final stage, stage four, begins above 2 000 °C with the annealing of the remaining distortions and the layers achieve flatness. At this point crystal growth can begin and graphitisation only truly starts. The thicknesses of the crystallites probably increase rapidly, accompanied by corresponding shrinkage in the other dimensions and the achievement of the final synthetic graphite microstructure.




Figure 2-30: Structural ordering during graphitisation (after Oberlin [23])

During graphitisation the crystallite sizes gradually increase from 5 nm to approximately 100 nm or more. This is readily confirmed by X-ray diffraction and high-resolution TEM analysis of materials at different stages of graphitisation

[24];[25]. However, industrial graphitisation does not produce a single graphite crystal. Hence macrostructures, e.g. optical texture, porosity, flow domains, etc., developed by the filler and binder phases remain after graphitisation. Ultimately these features, together with the extent of graphitisation, control the properties of the graphite [20].

The extent of graphitisation is a complex dependence on several steps, including defect removal from individual planes, linking of adjacent crystallites, evolution of the stacked structure, removal of dislocations and vacancies, etc. Thus at intermediate or even high levels of graphitisation the microstructure of the graphite can be a complex mixture of graphite crystallites of varying size and shape, interlinked by amorphous regions.

Understanding the oxidative behaviour of synthetic graphite materials is clearly highly dependent on the production method and raw materials used. Thus an extremely broad variety of conceivable materials result from this process and they have widely differing macro- and microstructures. Hence classifying and understanding the reactivity of these materials poses a significant challenge.

The presence of metallic catalysts dramatically influences the graphitisation [26];[27]. In general, however, due to the feedstock materials being processed hydrocarbons from petroleum refineries, synthetically derived graphite has very low impurity levels. However, graphite materials derived from other sources may have significantly differing microstructures. This is because graphitisation tends to be enhanced in the region of the catalyst particle. Near-perfect graphite crystallite formation around these particles leads to the development of onion-like structures [26]. The exact details of the morphology developed will depend on the precursor material and the catalyst in question.

2.3.4 Highly crystalline graphite

A special type of synthetic graphite known as “pyrolytic graphite” can be produced by chemical vapour deposition. This material is widely used as a coating on a variety of substrates. When suitable conditions and a hydrocarbon gas are used, the thermal decomposition of the gas leads to the deposition of columnar and laminar deposits on the substrate. In general, these deposits are turbostratic and graphitise readily when heat treated above 2 500 °C [22]. When

these deposits are annealed above 2 700 °C under a pressure of several atmospheres, a material known as “highly oriented pyrolytic graphite” (HOPG) is formed. This material is structurally very close to the ideal graphite crystal but is very expensive.

A related material is known as “Kish graphite”. This is a high-quality flake graphite produced during steelmaking. Due to the high carbon content of some steel melts, graphite will precipitate readily from these melts as they are cooled. The size and shape of these precipitates can differ widely [28], but in some cases large flat flakes in excess of a few millimetres wide and less than 100 µm thick may be obtained. These flakes have nearly perfect crystallinity and a low content of impurities.

2.3.5 Graphite characterisation techniques

Due to the extremely wide variety of possible graphite materials available, a large number of studies have focused on trying to characterise graphite samples from different origins [16];[29];[30]. However, the classification and techniques used vary widely, depending on the intended application of the graphite under consideration. Hence a single set of characterisation techniques with well-defined boundaries for classifying different graphite samples is not currently available. To limit the number of techniques covered, it is prudent to outline the pertinent properties that influence the oxidative behaviour of graphite.

The most relevant characteristic of graphite is its crystallinity. The extent of crystallinity in the two primary crystalline forms of carbon, namely diamond and graphite, is usually verified by a combination of three techniques [22]:

- A crystalline morphology visually observable by electron microscopy
- A single-phase crystalline structure detectable by X-ray or electron diffraction
- A clear peak on the Raman spectra, around 1330 cm⁻¹ for diamond and around 1580 cm⁻¹ for graphite

Other characteristics that are important for understanding the oxidation behaviour of graphite samples are: density, surface area, porosity, impurity content, reactivity, edge characteristics and particle size. The relevant analytical

techniques that allow characterisation of the previously mentioned properties will now be examined. The depth of the discussion will vary, depending on the level of expertise required to interpret the data correctly from the technique.

2.3.5.1 Macro- and microstructure characterisation – Microscopy

Optical microscopy offers a quick and effective technique for characterising the macrostructure and porosity of a sample. This technique uses visible light and lenses to illuminate and magnify a sample. However, it has several drawbacks, including the requirement for a flat surface which makes the analysis of powdered materials difficult. Furthermore, the resolving power of the microscope is limited by its construction and the wavelength of light. A variation on this technique is polarised light microscopy. As the name implies, this technique uses polarised light and allows the distinction between isotropic and anisotropic regions. It is very useful for distinguishing the anisotropic mesophase structures within a sample from the isotropic pitch regions.

Scanning electron microscopy (SEM) is widely used to characterise the microstructure of a wide variety of materials. This technique generates an image of the sample by scanning it with a high-energy beam of electrons in a raster pattern. The only requirement is that the sample must be electrically conductive. For non-conductive samples this limitation is usually overcome by sputter coating the sample with carbon or various metals, such as gold or platinum.

The electrons that are scattered by the sample are collected in a variety of ways. In general, two types of electron are detected: secondary electrons (SE) and back-scattered electrons (BSE). The detection rate of SEs depends on the orientation of the surface relative to the incident beam. Flat surfaces therefore tend to allow fewer electrons to be detected and appear dark, while steeply angled surfaces allow more, appearing brighter. For this reason SE detection is used mainly for topographical information.

Since heavy elements (high atomic number) backscatter electrons more strongly than light elements (low atomic number), and thus appear brighter in the image, BSE detectors are used to detect contrast between areas with different chemical compositions. Traditionally, SE detectors were the most widely used for generating images. However, due to recent advances in the field, including the

development of so-called “in-lens” detectors, the use of BSE detectors is increasing.

In a typical SEM, an electron beam is thermionically emitted from an electron gun fitted with a tungsten filament cathode. The resolution of a modern SEM has been further enhanced by the increasing use of a field emission gun as the electron source. These guns produce an electron beam that is smaller in diameter, more coherent and has greater current density than the older type. These advancements in the field of electron microscopy allow unprecedented resolution (typically <1 nm) at very low voltages. Operation at low voltages limits the penetration of electrons into the sample, thus ensuring that only the top layer of material is analysed. This allows excellent resolution of surface detail and structure.

An alternative to SEM imaging is transmission electron microscopy (TEM). Electrons are utilised as in the SEM but in this case they are transmitted through the sample, interacting with the sample as they pass through. They are then collected and an image is constructed. The primary requirement for TEM analysis is an ultra-thin sample to allow sufficient electrons to pass through. This makes the technique unsuitable for bulk or powder analysis. However, atomic resolution is possible in modern TEMs, making them useful for studying the edges of single, or sometimes multilayered, graphene samples [14]. A related technique, scanning tunnelling electron microscopy (STEM or STM), has recently [31] shown the potential to identify oxygen edge groups on graphite, but more work is needed to gain more useful insights.

2.3.5.2 Crystal characterisation – XRD

X-ray diffraction (XRD) is a method for determining the structure of atoms within a crystal lattice. A beam of X-rays strikes a crystal and the rays are diffracted into a variety of specific directions, as shown in Figure 2-31. From the angles and intensities of these beams the three-dimensional structure of the crystal may be inferred. This is done by the use of, among other relationships, Bragg’s law which may be briefly described as follows.




Figure 2-31: Electromagnetic beam incident on a crystal lattice

Crystals are composed of regular arrays of atoms and X-rays are considered as electromagnetic waves. When the X-rays strike the atoms they are scattered primarily by the atoms' electrons via elastic scattering. In most directions these scattered waves cancel each other out via destructive interference. However, in certain directions these waves undergo constructive interference, as determined by Bragg's law:

$$2d' \sin \theta = n\lambda \quad (2.3)$$

Here d' is the spacing between diffracting planes, θ is the angle of the incident beam, n is any integer and λ is the wavelength of the beam.

This technique has been used to characterise carbon materials since the early 19th century when Bragg used it to identify the crystal structure of diamond [32]. Shortly after, Debye and Scherrer [33] solved the structure of graphite by the related method of powder diffraction. The wavelengths of some commonly used radiation sources are shown in Table 2-8.

Table 2-8: X-ray spectral line wavelengths (nm) for common sources [34]

Source	$K\alpha$ (nm)
Fe	0.19373
Co	0.17903
Ni	0.16579
Cu	0.15418
Zr	0.07859
Mo	0.07107

During powder X-ray diffractometry (PXRD) it is assumed that every possible crystal orientation is represented equally in the powdered sample. The resulting orientational averaging causes the three-dimensional reciprocal space (Fourier transformed space) that is normally studied in single crystal diffraction to be projected onto a single dimension. In this case the intensity should be homogenous across the sample. When the scattered radiation is collected on a flat plate detector, the rotational averaging leads to smooth diffraction rings around the beam axis. The angle between the beam axis and the ring is called the “scattering angle” and is always denoted as 2θ .

Modern powder diffraction data are usually represented as a diffractogram in which the intensity is shown as a function of the scattering angle, as opposed to the raw diffraction ring data. The position of a diffraction peak is independent of the atomic positions within the cell and entirely determined by the size and shape of the unit cell of the specific crystalline phase. Hence each peak represents a certain lattice plane and can therefore be characterised by a Miller index.

The ideal XRD stick spectra for hexagonal and rhombohedral graphite materials are shown in Figure 2-32 and Figure 2-33 respectively, complete with corresponding Miller indices using $\text{CoK}\alpha$ as the radiation source. These spectra were obtained from the ICDD database (PCPDFWIN Version 2.4) used in the PANalytical X'Pert Pro powder diffractometer for $\text{CoK}\alpha$ radiation. The intensity is on a log scale to allow better visualisation.

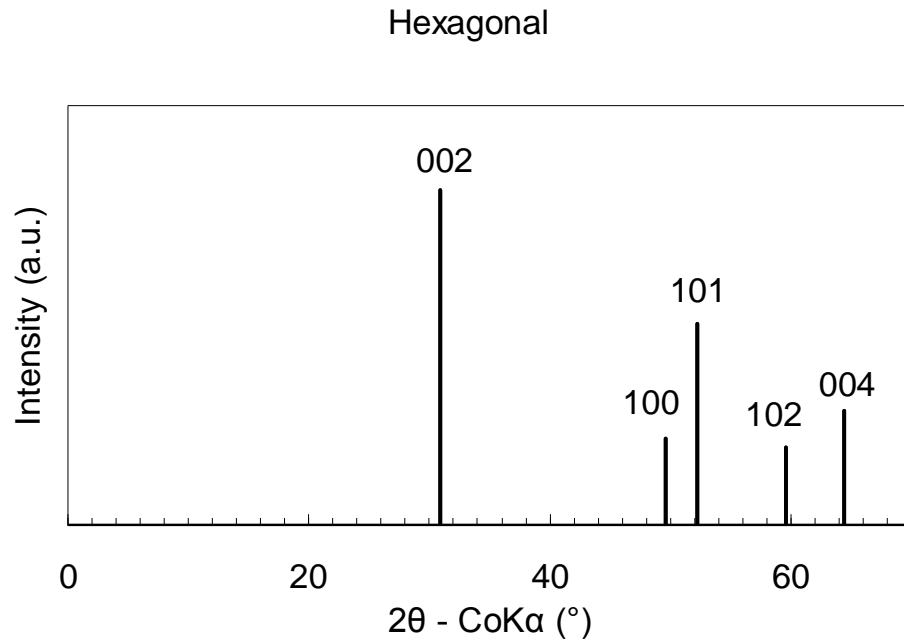


Figure 2-32: XRD stick spectrum of ideal hexagonal graphite crystal




Figure 2-33: XRD stick spectrum of ideal rhombohedral graphite crystal

The most dominant line in the XRD spectra for hexagonal graphite is the $\langle 002 \rangle$ peak. This peak can be used to determine the interlayer spacing (d'_{002}) of graphite using Eq. (2.3). Since the peak occurs at $2\theta = 30.92^\circ$, the corresponding interlayer spacing for ideal hexagonal graphite is $d'_{002} = 0.3354$ nm. For PXRD,

the intensity should be independent of all other aspects, except the angle of the incident beam. Hence any line-broadening effects for a pure material must be caused either by instrument factors or by crystallographic effects. These effects include the presence of defects in the lattice, strain differences in different grains or differences in crystallite sizes. Some of the first work on this phenomenon was done by Scherrer [35] who formulated the empirical Scherrer equation:

$$L = \frac{K^h \lambda}{\beta^h \cos \theta} \quad (2.4)$$

where θ is still the angle of the incident beam, λ is the wavelength of the beam, L is the size of the crystalline domains, K^h is the shape factor and β^h is the peak width at half maximum intensity.

This work showed [36] that for the determination of crystallite height due to crystalline reflections perpendicular to the ideal crystal plane, e.g. {002}, or rather along the c-direction of the unit cell, i.e. L_c , the shape factor is equal to 0.9 and the $\langle 002 \rangle$ peak may be used for the peak width calculation.

One of the first studies to characterise peak broadening for carbon materials, namely heat-treated carbon black, was done by Warren in 1941 [37]. He showed that for two-dimensional crystallite size determination along the ab-direction of the unit cell, i.e. L_a , the shape factor is equal to 1.84 when determined from a given $\langle hk \rangle$ line. Here L_a is considered to be the diameter of the crystallite and usually either the $\langle 101 \rangle$ ($2\theta = 52.2^\circ$, $\text{CoK}\alpha$) or $\langle 110 \rangle$ ($2\theta = 93.1^\circ$, $\text{CoK}\alpha$) peak is used for its calculation.

These concepts are neatly illustrated by a model developed by Fujimoto and Shiraishi [38]. The graphite crystallites are modelled using a multilayered, quasi-circular arrangement of hexagonal rings of increasing size, as shown schematically in Figure 2-34. This approach was successfully applied to mesocarbon microbeads to analyse the crystallite development as a function of heat-treatment temperature.




Figure 2-34: Graphite crystal model

In recent times new work has been done to further enhance the value gained from the XRD spectrum. Lu *et al.* [39] developed a new approach based on a reduced intensity. This intensity is derived from an intensity curve normalised to electron units and the theoretical curves for the coherent and incoherent scattering of the carbon atom. Based on this, the fraction of amorphous material, and the crystallite height and diameter can be calculated. This was done for a variety of Australian black coals. The results show good agreement with the Scherrer equation. Furthermore, the aromaticity of the coal is inferred via the asymmetry of the $\langle 002 \rangle$ peak. This aromaticity shows good correlation with values measured for other coals using nuclear magnetic resonance (NMR) imaging, as well the hydrogen-to-carbon ratio of the coals, thus illustrating the wide spread of information that can be obtained from XRD.

Several studies [40];[41];[42] have focused on studying the stacking disorder in graphite resulting from a mixture of hexagonal and rhombohedral phases. This has been done with models of varying complexity and with varying success. The rhombohedral and hexagonal lattices have slightly differing edge structures and therefore may have different reactivities. Thus it is useful to quantify the relative amounts of these phases in a sample. The simplest approach is given by Parthasarathy *et al.* [8] by taking the intensity ratios of the $\langle 101 \rangle$ diffraction peaks for each of the phases.

A more thorough approach was developed by Shi *et al.* [43]. Using a two- and three-layer graphite structural model, a curve fit for the raw XRD intensity data is calculated. The structural parameters of the model are then refined to achieve an optimal fit. Based on these parameters, the rhombohedral content and the stacking disorder in the rhombohedral and hexagonal phases can be

calculated. The model was successfully applied to determine the lithium intercalation properties of the graphite.

2.3.5.3 Atomic and crystal structure – Raman, Nuclear magnetic resonance

Raman scattering is the inelastic scattering of photons by phonons due to the change of polarisation caused by the phonon mode [44]. It is used to study the vibrational, rotational and other low-frequency modes in a system. In short, when a photon impinges upon a molecule, it excites the molecule from its ground state. When the molecule relaxes back down, it emits a photon and returns to a new vibrational or rotational state. The difference in energy between the initial and final states causes a shift in the frequency of the emitted photon away from the excitation wavelength. This difference in wavelength is known as the Raman shift wave number.

As mentioned earlier, single-crystal graphite belongs to the d_{6h}^4 symmetry group and therefore has vibrational modes, $2E_{2g}$, $2B_{2g}$, E_{1u} and A_{2u} , of which only the $2E_{2g}$ modes are Raman active [45];[46]. This mode is active at a Raman shift wavenumber of 1580 cm^{-1} and is generally known as the G-band. Polycrystalline graphite shows an additional peak around 1360 cm^{-1} , which is generally known as the D-band. Tuinstra and Koenig [47] attributed this band to an A_{1g} mode which becomes active due to the finite crystal size. A schematic representation of both modes is shown in Figure 2-35.




Figure 2-35: Raman active modes for graphite

Using this assumption, they correlated the ratio of the absolute D- to G-band intensities, i.e. $I(D)/I(G)$, to the crystallite radius, L_a (measured by XRD) for a variety of carbons [47], yielding the empirical relationship:

$$\frac{I(D)}{I(G)} = \frac{C(\lambda)}{L_a} \quad (2.5)$$

where C is a correlation coefficient that is dependent on the Raman wavelength used, due to the dependence of the D-band on wavelength.

For their work Tuinstra and Koenig found $C(515.5 \text{ nm}) \sim 4.4 \text{ nm}$. Nikiel and Jagodinski [48] repeated this work for a variety of nuclear-grade graphite samples, HOPG and “Grafoil”. They obtained a similar linear relationship but with a vastly different correlation coefficient.

Wang *et al.* [46] point out that a small L_a is a sufficient but not necessary condition for the appearance of the D-band. They indicate that the D-band is affected by the presence of edges in the material which are not necessarily the consequence of decreased crystallite size. This is done by showing the presence of the D-peak in HOPG material which is analysed with incident radiation polarised parallel to the c-axis, as well as being present in HOPG material that is electrically oxidised along channels in the basal plane.

This conclusion is further supported by Kawashima and Katagari [45] who conducted similar analyses on HOPG and pyrolytic graphite. They raise further questions surrounding the calculation of the existing density of states to explain the phonon spectrum of graphite, stressing that additional work is needed. Thus the utility of the Raman spectra to infer the size of crystallites in graphite samples seems questionable. However, it has a possible application in discerning between sp^2 and sp^3 hybridised materials and hence potentially determining the amorphous content of a sample. This particular aspect will be discussed in more detail in Section 2.4.

Further work done by Nikiel and Jagodinski [49] indicates that the Raman spectra of nuclear-grade graphite samples are exceptionally sensitive to species adsorbed on the graphite surface. It appears that the changes in the Raman spectra are related to the trace elements that are present in the graphite samples. Some of the adsorbed material could be removed by heating under vacuum. This study indicates an additional difficulty when interpreting the Raman spectra of samples with unknown histories and impurity contents.

A technique less widely used to characterise the atomic structure of carbon materials is nuclear magnetic resonance (NMR). NMR spectroscopy has been successfully applied to a wide variety of carbon materials [50] and has even been used to monitor graphitic structural development during systematic heat treatment [51]. However, highly ordered graphite materials have a highly anisotropic magnetic susceptibility and disordered carbon materials exhibit high chemical shift anisotropy. Due to these facts, the NMR spectra of carbon materials are usually broad and have poor resolution. High-resolution techniques do exist but the high electrical conductivity of graphite causes tuning problems. As such, these techniques can be cumbersome and time consuming to use and obtain useful data.

2.3.5.4 Density, porosity and surface area – Pycnometry, mercury intrusion and gas adsorption

For bulk materials or even powdered samples derived from bulk materials, understanding the macrostructure is very important from an oxidation perspective. Microscopy allows only the outer edge of the sample to be visually inspected for its macro- and microstructure, with the possible exception of the recently developed focused ion beam technology. In order to quantify the structural parameters of the sample and gain an insight into the internal structure, alternative techniques are required. The simplest method for determining the density of an irregularly-shaped sample is to submerge it in a liquid and measure the volume of fluid displaced. This method is known as pycnometry and is suitable for non-porous and some porous materials. In modern pycnometers the displaced fluid may be a gas, liquid or very fine powder.

Depending on the porosity of the material, the definition of density may vary. The skeletal density (also termed the true, real, apparent or absolute density) is obtained when the volume measured excludes the pores and void spaces in between particles in a bulk sample. To ensure that all the pores are penetrated by the fluid, a gas, mostly helium, is used for this measurement, leading to the additional term “helium density”. The bulk density (or envelope density) is the term used for porous materials when the pore spaces within the particles are included in the volume measurement. Thus it is critical that the fluid

utilised in the measurement does not penetrate the pores. For this reason, a free-flowing dry powder is usually used as the displaced medium. For non-porous materials the two types of density are equivalent.

A third definition of density, known as the tap density, is also possible for powdered or particulate samples. In this case the bulk density of the sample is determined after the sample has been vibrated to obtain, as near as possible, the optimal packing. These techniques will allow the total porosity or void fraction to be determined. However, no information regarding the size, shape or distribution of the pores is obtained.

To determine the pore size distribution, mercury intrusion porosimetry is often used. Since mercury does not wet most substances and will not spontaneously penetrate the pores by capillary action, it must be forced into the pores by the application of external pressure. The required equilibrated pressure is inversely proportional to the size of the pores, only slight pressure being required to intrude mercury into large macropores, whereas much greater pressures are required to force mercury into small pores.

Mercury porosimetry analysis is the progressive intrusion of mercury into a porous structure under stringently controlled pressures. From the pressure versus intrusion data, the instrument generates volume and size distributions using the Washburn equation. Clearly, the more accurate the pressure measurements, the more accurate the resulting pore size data. Due to the high pressures, it is possible that this technique may damage the pore structure of soft materials such as graphite.

Synthetic graphite is a particularly problematic material for pore size and distribution determination since the material contains a uniquely complex pore structure. Firstly, a significant percentage of the porosity of the material is contained in pores that are not open [52]. This closed-pore volume (CPV) cannot be accessed by any of the intrusion-based methods mentioned previously. This limitation can be overcome to a certain extent, in terms of the void fraction determination, if the skeletal density can be independently determined or verified to be equal to the ideal crystal density of graphite, within some reasonable approximation.

Secondly, these materials also contain so-called “restricted access pores” (RAPs). These “ink-bottle”-shaped pores [52] have a small inlet aperture which

leads to a much larger open cavity. This significantly distorts the interpretation of mercury porosimetry data since the volume associated with the small pores will be significantly overestimated and clear hysteresis will be noticeable in the intrusion curves.

These limitations can be overcome by using a technique known as gas adsorption. In this context adsorption is the adhesion of gas molecules to a surface, thereby creating a film of the adsorbate on the surface of the adsorbent. This technique generally applies to physisorption, whereby the adsorbate is held in place by physical forces such as weak van der Waals or electrostatic interactions. An alternative is chemisorption, whereby the adsorbate is attached to the adsorbent via a chemical bond. This technique will be discussed in more detail in Section 2.7.

Most commonly, the technique of Brunauer, Emmett and Teller [53] is used to determine the so-called “BET surface area”. This method relies on the assumptions that the gas adsorbs onto the solid in layers *ad infinitum*, there is no interaction between adjacent layers and the gas coverage can be described by the Langmuir expression. At least three points in the relative pressure range, i.e. pressure divided by vapour pressure, are chosen, usually between 0.05 and 0.35. The weights of gas adsorbed at these pressures are measured and plotted using the BET equation. Using the molecular cross-sectional area of the gas adsorbed, the monolayer coverage can be estimated.

In general, a pore size distribution is described through isotherms, i.e. the amount of adsorbate on the adsorbent as a function of the gas pressure. However, this technique has some pitfalls and limitations [54]. In order to ascertain the pore size distribution from the adsorption isotherms, a variety of mathematical models are used. These models depend heavily on an a priori knowledge of the microstructure of the material. As mentioned earlier, graphite and, even more so, amorphous carbon materials [53]. S. Brunauer, P. H. Emmett and E. Teller (1938) Adsorption of Gases in Multimolecular Layers *Journal of American Chemical Society*, **60**, 309–319.

[54] J.C. Groen, L.A.A. Peffer and J. Pérez-Ramírez (2003) Pore size determination in modified micro- and mesoporous materials. Pitfalls and limitations in gas adsorption data analysis. *Microporous and Mesoporous Materials*, **60**, 1–17.

[55] have very unique and complex microstructures. This makes correct interpretation of the adsorption isotherms for these materials problematic, requiring increasingly complex modelling approaches [56];[57].

Recently, excellent work has been done using X-ray tomography to characterise the pore network of nuclear graphite [58];[59]. However, currently this technique is being utilised only for macroscopic studies with a maximum voxel resolution of $13 \mu\text{m}^3$. If the technique could be refined to improve the resolution, it may have future use in thoroughly characterising the pore structure of graphite and related carbon materials.

2.3.5.5 Impurity content – Compositional spectroscopy

There is a wide variety of spectroscopic techniques for compositional analysis. One of the most common techniques for analysing trace elements is inductively coupled plasma–mass spectrometry (ICP-MS). Using an extremely high temperature ($>10\,000\text{ K}$), samples of argon plasma, which is generated with a high-frequency induction coil, are ionised into individual atoms. Using a mass spectrometer (MS), usually a quadrupole, the ions are separated on the basis of their mass-to-charge ratio and a detector receives an ion signal proportional to the concentration.

The absolute concentration of a sample can be determined through calibration with certified reference materials, such as single- or multi-element reference standards. A related technique, ICP-AES, uses atomic emission spectroscopy (AES) to identify the elements. Both techniques suffer the significant drawback that the sample must be completely dissolved in a suitable liquid medium. In most cases samples are ashed and digested by a mixture of acids. Alternative methods for sample introduction, such as laser ablation and electrothermal vaporisation, are available but their use is not widespread.

A technique that is very suitable for bulk sample analysis is X-ray fluorescence (XRF) spectroscopy. This technique utilises the fact that materials emit characteristic "secondary" (or fluorescent) X-rays when they are bombarded with high-energy X-rays. Sample penetration is fairly limited ($<100 \mu\text{m}$), depending on the energy, or wavelength, of the incident X-rays. For this reason the sample should be very homogenous to give an accurate representation of the

overall composition. Depending on the elements being analysed and the instrumental set-up, detection limits of a few ppm can be readily achieved.

At rest, an atom within the sample contains ground state (or unexcited) electrons in discrete energy levels or electron shells bound to the nucleus. The incident beam may excite an electron in an inner shell, ejecting it from the shell while creating an electron hole where the electron was. An electron from an outer higher-energy shell then fills the hole, and the difference in energy between the higher-energy shell and the lower-energy shell may be released in the form of an X-ray, as shown schematically in Figure 2-36.




Figure 2-36: Atomic X-ray emission

For example, if an electron from the L-shell fills a hole in the K-shell, a $K\alpha$ X-ray will be emitted. Similarly, if an electron from the M-shell fills a hole in the K-shell, a $L\alpha$ X-ray is emitted. For any given element there is a variety of possible emissions, giving rise to a spectrum specific to that element. In theory, any emission line can be used to identify a given element since, for example, the $K\alpha$ line is unique for each element. However, it is possible that X-ray emission lines from certain elements overlap with those from other elements, but from other shells. Such overlaps should be borne in mind and compensated for, especially during quantitative analysis. In general, this is achieved through complicated, proprietary routines in the identification software.

Detection of the emitted X-rays can be achieved in a variety of ways, depending on the application. The first technique is so-called energy dispersive X-ray (EDS or EDX) spectroscopy. In energy dispersive analysis, dispersion and

detection are a single operation and various types of solid state detector are used. A single incoming X-ray photon ionises a large number of detector atoms, with the amount of charge produced being proportional to the energy of the incoming photon. The charge is then collected and the process repeats itself for the next photon. An example of a spectrum generated in such a fashion is shown in Figure 2-37.

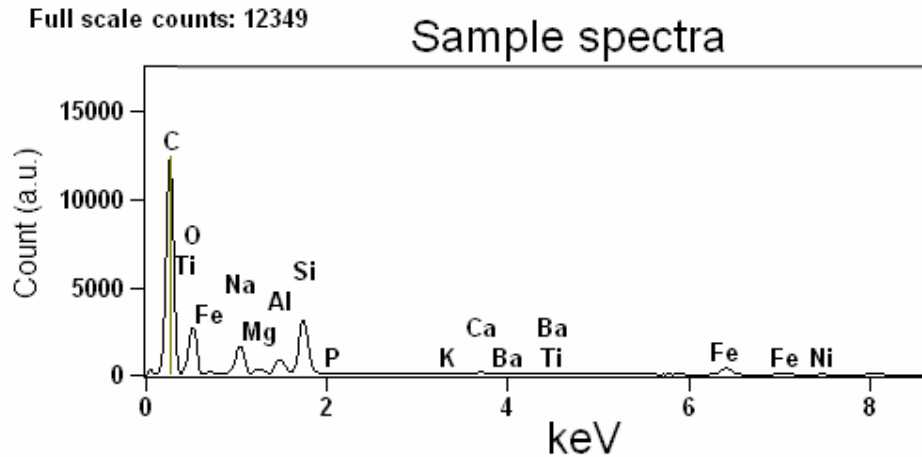


Figure 2-37: Example of an EDS spectrum

The alternative technique is wavelength dispersive X-ray (WDS or WDX) spectroscopy. In wavelength dispersive analysis, the fluorescent X-rays emitted by the sample are directed into a diffraction grating monochromator. The diffraction grating used is usually a single crystal. By varying the angle of incidence and take-off on the crystal, a single X-ray wavelength can be selected. Both techniques have different advantages and drawbacks. For example, EDS allows rapid acquisition of a full spectral analysis, whereas WDS allows much higher sensitivity for specific wavelengths, making quantitative analysis more accurate.

Since the emission of secondary X-rays from a specimen can be induced by a high-energy beam of X-rays or charged particles, such as electrons or protons, these techniques are often coupled with SEM. However, this adds complexity in terms of the beam current and acceleration voltage required in the SEM to induce sufficient X-ray emission for analysis. Furthermore, the resolution of this technique is limited by the size of the incident electron beam. Hence it is

not possible to resolve the composition of small particles on the surface of the sample. Instead, the analysis will contain X-rays from the entire region affected by the electron beam. This is further complicated by uneven samples, in which case the surface geometry and path of X-rays from the sample to the detector should be considered.

A technique that has recently gained increasing popularity is electron energy loss spectroscopy (EELS). This technique is used in conjunction with TEM analysis. The electrons that are transmitted through the sample are analysed for their energy loss due to inelastic scattering. Inelastic interactions include phonon excitations, inter- and intra-band transitions, plasmon excitations, inner-shell ionisations and Čerenkov radiation. The inner-shell ionisations are particularly useful for detecting the elemental components of a material. Recent progress in techniques for energy-based filtering of the transmitted electrons has dramatically increased the energy resolution of this technique. Currently, atomic resolution of composition is possible, depending on the specific sample and conditions. However, since the electrons must be transmitted through the sample as in TEM, this technique is limited to ultra-thin samples only.

2.3.5.6 Reactivity – Thermal analysis

Several thermal analysis techniques are readily available for characterising samples as a function of temperature. The technique most relevant to reactivity analysis is thermogravimetric analysis (TGA). The analyser usually consists of a high-precision mass balance with a sample container. A common design utilises a balance beam with a specifically calibrated counterweight. Usually, a secondary reference beam, which contains no sample, is used to compensate for potential measurement errors such as beam growth and buoyancy effects. Embedded within the beams and located beneath the sample holder are thermocouples, which are used to determine the sample temperature accurately. A schematic representation of this set-up is shown in Figure 2-38.




Figure 2-38: Schematic of TGA

The entire beam construction is placed within a small electrically heated oven. The atmosphere in the reaction chamber may be continuously purged with gas to maintain an inert or oxidative atmosphere. The entire instrument is computer controlled and weight, temperature, heat input and purge gas flow are continuously monitored. Under inert conditions the system can be used to identify moisture and volatile content. This ability is further enhanced by the simultaneous analysis of the exhaust gas using a mass spectrometer (MS), which allows continuous monitoring of gas composition to identify the released species.

Under oxidative conditions the mass measurement can provide insight into the sample reactivity at a variety of isothermal and non-isothermal conditions. A variety of temperature programs and conditions can usually be selected, depending on the desired result. Since the mass measurement is absolute, this technique provides the ash content of the sample if it is oxidised to full conversion. Thermal analysis data for solid reactants are frequently analysed using the solid state kinetic approach which is discussed in Section 2.12.

2.3.5.7 Edge characterisation – Various techniques

The correct nature of the plane edges of graphite is still not exactly known [60]. It is widely accepted that a variety of surface oxygen complexes are chemically bound to these edges. Over the years several techniques have been

utilised to analyse these edges [61]. A technique that is related to XRF is X-ray photoelectron (XPS) spectroscopy. In this case, however, electrons emitted via the photoelectric effect are detected when the sample is irradiated with X-rays. This technique has been successfully applied to study the chemisorption of oxygen on the basal plane of graphite [62].

The amount of electrons emitted is not only a function of the incident radiation, but also the mean free path to escape. Since these are low-energy electrons, they are generally recaptured if the depth below the surface is significant. Typically, this technique analyses only the upper 10 nm of the sample surface and therefore detection limits are in the order of parts per thousand. This technique also requires ultra-high vacuum conditions to prevent scattering of the electrons.

A technique similar to Raman spectroscopy is Fourier transform infrared (FTIR) spectroscopy. This technique utilises electromagnetic radiation from the infrared part of the spectrum, as opposed to Raman spectroscopy which generally utilises the visible or near-ultraviolet region. However, in this case the aspect of interest is the absorbance of the electromagnetic radiation rather than a shift in its wavelength. Light from a polychromatic infrared source is collimated and directed at a beam splitter. Ideally, half the light is reflected towards a fixed mirror and the other half is directed through the sample at a moving mirror.

The reflected light from the two mirrors is recombined and directed towards a detector. By varying the distance of the moving mirror, an interferogram can be constructed. The interferogram is based in the length domain; a Fourier transform is used to invert into the wavenumber domain. By using this technique the absorption spectra of the sample can be obtained across a wide range of wavenumbers, i.e. the full spectrum of wavelengths present in the original polychromatic source.

The absorption of light is dependent on the specific chemical composition of the sample. In general, each compound has a unique spectral response. Thus its presence in the sample can be identified by comparing the spectrum obtained against a database containing a collection of likely candidates. In addition, the absorption spectrum is sensitive to certain types of chemical bond. This allows the identification of different oxygen-based structures within the sample. For example, the presence of carbonyl groups, i.e. carbon double bonded to an

oxygen, may be detected by strong absorbance in the region of 1700 cm^{-1} . Similarly, a carbon-oxygen single bond, as found in ether structures, absorbs in the region of $1000\text{--}1300\text{ cm}^{-1}$.

There is a variety of infrared techniques, including direct transmission and attenuated total reflectance (ATR), whereby the infrared beam samples a surface by internal reflection of the light. Diffuse reflectance infra-red Fourier transform (DRIFT) occurs when light impinges on the surface of a material and is partially reflected and transmitted. Light that passes into the material may be absorbed or reflected out again. Hence the radiation that reflects from an absorbing material is composed of surface-reflected and bulk re-emitted components, which summed are the diffuse reflectance of the sample

The simplest technique is related to thermal analysis: temperature programmed desorption (TPD). In this technique a sample is systematically heated in a TGA which is coupled to an MS. As the surface groups are thermally desorbed, they register as either carbon monoxide or carbon dioxide in the MS. Unfortunately, this method must be used in conjunction with other methods to discern the bonding structure of the oxygen complexes that were desorbed within a specific temperature range.

A technique that has been applied less than the others is acid-base titration [63];[64]. This technique uses potentiometric titration to analyse surface groups on carbon materials. It has been successfully applied to activated carbons to identify surface groups with acid and basic characteristics. However, it provides no way of directly determining the bonding structure or composition of these groups. It can only be assumed that they are carboxylic acid or hydroxyl based. Furthermore, constrictions in the microporosity and hydrolysis of the chemical groups may lead to incorrect estimates of the total surface chemical groups by this method [65].

The application of all of the surface characterisation techniques to graphite and other carbon materials will be discussed in more detail in Section 2.8.

2.3.5.8 Particle sizing – Laser diffraction

For powdered samples accurate knowledge of the particle size distribution is of paramount importance. There is a variety of techniques for analysing the

particle size, the most widely utilised being dynamic light scattering. With this technique a focused laser beam is directed at a fluid cell in which the particles are dispersed. The particles scatter the light at an angle which is inversely proportional to their size. A map of the scattering intensity against the scattering angle is used to infer the particle size distribution from Mie theory.

This technique has several complications, including the need for adequate sample dispersion in the fluid medium, i.e. very little settling, as well as the underlying Mie theory assumption that the particles are spherical in shape. These factors may lead to erroneous estimation of the particle size. Graphite powders have the added disadvantage that they tend to agglomerate, thus making the choice of dispersant even more critical. Alternatively, traditional sieving may be used to determine the particle size distribution.

2.3.6 Graphite properties [22]

A few of the physical properties for the theoretically ideal graphite crystal are shown in Table 2-9. Graphite has the highest melting point of any element and the only better refractories are hafnium carbide and tantalum carbide.

Table 2-9: Physical properties of graphite

Density (300 K, 1 atm)	2.26	g/cm ³
Atomic volume	5.315	cm ³ /mol
Triple point (estimated)	4 200	K
Boiling point (estimated)	4 560	K
Sublimation point at 1 atm (estimated)	4 000	K
Heat of fusion	46.84	kJ/mol
Pauling electronegativity	2.5	

The thermal properties of graphite are given in Table 2-10. Note that the unique crystal structure of graphite gives rise to considerable anisotropy in the properties when measured along the graphite plane (*ab* direction) or perpendicular to it (*c* direction).

Table 2-10: Thermal properties of graphite

Heat of combustion ΔH_{CO} @ 25 °C and constant pressure to form CO_2 gas	393.13	kJ/mol
Standard entropy S° @ 25 °C	5.697–5.743	J/mol.K
Entropy ΔS_{298}	152.3	kJ/mol
Enthalpy ΔH_{298}	716.88	kJ/kg.K
Specific heat @ 25 °C	0.690–0.719	W/m.K
Thermal conductivity: <i>ab</i> direction @ 25°C	398	W/m.K
Thermal conductivity: <i>c</i> direction @ 25°C	2.2	W/m.K
Coeff. of thermal expansion: <i>ab</i> direction (100 – 600 °C)	$3.2\text{--}7.0 \times 10^{-6}$	m/m.°C
Coeff. of thermal expansion: <i>c</i> direction (100 – 600 °C)	25×10^{-6}	m/m.°C
Resistivity: <i>ab</i> direction @ 25 °C	$2.5\text{--}5.0 \times 10^{-6}$	ohm.m
Resistivity: <i>c</i> direction @ 25 °C	$3\,000 \times 10^{-6}$	ohm.m
Young's modulus: <i>ab</i> direction	1 060	GPa
Young's modulus: <i>c</i> direction	36.5	GPa
Young's modulus: parallel to plane	4.5	GPa

Some of the properties given in this section depend on the size and orientation of the crystallites within the graphite. As a result, a range for the property value is given in some cases and, where possible, the value for the most crystalline form is given. The values are meant as indicators or for comparison only and are not representative of any particular graphite.

The thermal conductivity is that of pyrolytic graphite and is a strong function of temperature. The coefficient of thermal expansion in the *ab* direction is that of bulk graphite and is also a strong function of temperature. Due to its planar nature, graphite has a very low shear resistance, hence the low value for the Young's modulus parallel to the plane, or shear modulus.

2.4 Amorphous carbon

Between the perfectly crystalline structure of the two main carbon allotropes, graphite and diamond, one finds a limitless myriad of possible carbonaceous materials. Broadly, these carbons are classified as amorphous due to the lack of long-range structural order. In general, they include anything from coal to activated carbon. Their actual structure is still poorly understood despite numerous investigations into a wide variety of these materials [39]; [66].

Technically, some of these materials are not strictly speaking pure carbons, as they naturally contain significant amounts of elements other than carbon. Since many of these materials are derived from hydrocarbons, they typically contain a significant amount of hydrogen. These can be present as volatile (mostly organic) compounds which are driven off by heating to medium temperatures. Or they may be present via direct bonding to the carbon substructure, which requires much higher temperatures to remove in a step known as carbonisation. However, the combustion of these materials has historically been considered as part of the broader investigation into the reaction of carbon with oxygen. This is especially true of the forms of these carbons that have been carbonised at high temperature, such as chars and cokes. These differences should, however, be borne in mind when reading Section 2.5 in which the oxidation behaviour of carbon materials in general is considered.

These amorphous carbons are presumed to be less crystalline varieties of graphite or diamond, or some combination thereof. However, the exact atomic structure of these materials still remains somewhat of a mystery. Raman spectroscopy appears to be the ideal method for characterising materials that are intermediate between graphite and diamond. Both perfectly crystalline graphite and diamond have clear, distinct Raman patterns, as shown in Figure 2-39. Graphite shows a narrow, tall peak around 1570 cm^{-1} , sometimes referred to as the G-band, while diamond shows a tall narrow peak around 1330 cm^{-1} , sometimes referred to as the D-band. This is based on work done on either highly ordered pyrolytic and natural graphite samples or gem-quality natural diamonds [67].




Figure 2-39: Raman spectra of diamond and graphite (from Pierson [22])

It may be inferred that the G-peak is an indicator of the sp^2 -hybridised material within a sample, while the D-peak is an indicator of the amount of sp^3 material in the sample. Hence for amorphous material it would be tempting to assume that a smooth transition occurs, leading to some intermediate spectrum which is a predictable combination of these two spectra. However, the situation is slightly more complex than this.

In order to build an appreciation for the spread of structures that is possible in amorphous materials, it is perhaps best to start with the transition that graphite undergoes as it is progressively milled. Welham *et al.* [68] extensively ball milled a high-purity graphite sample (>99.9% C). As mentioned in Section 2.3.5.3, the ratio of the D- and G-peak intensities is related to the size of crystallites in the graphite material. Indeed, they found that progressive milling caused a dramatic increase in the height of the D-peak relative to that of the G-peak. This indicates that a decrease in the crystallinity of the graphite had occurred. Earlier work by these researchers [69] reinforces this conclusion since the XRD intensity of the $\langle 002 \rangle$ peak of ball-milled graphite rapidly decreased, indicating the destruction of crystalline ordering and the formation of amorphous carbon. This was further confirmed by comparing the XRD trace with that of activated carbon, which is known to have an amorphous structure. Both traces showed clear similarity and shared trace structures.

Furthermore, Welham *et al.* proceeded to fit five Lorentzian peaks at wavenumbers 1166, 1309, 1473, 1573 and 1607 cm^{-1} respectively, as shown in Figure 2-40. They state that the peak at 1166 cm^{-1} is evidence of sp^3 -hybridisation in the molecular structure. However, the relative intensities of all peaks, except the D-peak, are very low with an appreciable amount of noise on the fitted spectrum. This leads to some ambiguity regarding the conclusions. The work of Endo *et al.* [70] explicitly state that the appearance of the D-band indicates the presence of disorder within the graphitic structure, but do not link this to the presence of sp^3 -hybridised material. The statement by Welham and co-workers that the peak found around 1150 cm^{-1} is related to sp^3 content is justified by three references.




Figure 2-40: Raman spectra of milled graphite (from Welham *et al.* [68])

In the first of these [71], transmission electron energy loss spectroscopy (EELS) is used to determine the sp^3 content of a variety of diamond-like carbon films. This technique is not suitable for bulk characterisation as only the top few layers can be analysed due to the low probe depth. However, thin films were used in this study and it was found that above approximately 30% sp^3 -bonded carbon atoms, a persistent peak develops at a wavenumber of 1150 cm^{-1} . It is highly unlikely that such a high sp^3 content could have been achieved via milling alone. Below this threshold value, the peak shifts towards 1300 cm^{-1} and can only readily be seen by fitting and subtracting a curve to the G-band. Furthermore, no reference is made to the D-band and it may well be that this new peak is a result of a shifting in the traditional D-band due to effects during the formation of the film under consideration.

The second reference [72] refers to a tight-binding molecular-dynamics simulation where a hypothetical three-dimensional lattice of carbon is constructed with a varying amount of atoms with three and four bonds. The vibrational properties of this hypothetical structure are analysed. It is concluded that a broad peak between 800 and 1000 cm^{-1} may be the result of three- and four-fold atomic interactions, or regions where the four-fold bond angles are severely distorted. This conclusion is based on a model with, again, a very high percentage (74%) of

sp^3 -bonded material and the resulting peak is far lower than peak under consideration.

The final reference [73] contains no information on a peak around 1150 cm^{-1} and instead refers to the development of a peak around 1350 cm^{-1} due to nitrogen doping. It is interesting to note that this new peak is absent from all but one of the spectra of amorphous and diamond-like carbon films characterised by Knight and White [67]. Thus it is evident that there is no clear link to sp^3 content from the peak at 1150 cm^{-1} [74], nor from any other peak in this work. The result of this milling experiment potentially indicates simply that sp^2 crystallites have been decreased in size.

Next, one might consider the possible derivatives of diamond with a damaged crystal structure, such as so-called diamond-like carbon (DLC) or tetrahedral amorphous carbon (*ta-C*). DLC can be considered a metastable form of carbon which consists of a predominantly amorphous structure and a significant fraction of sp^3 bonds. The molecular structure may be in the form of individual or polycrystalline diamond, in small clusters (5–10nm), interspersed with amorphous regions [22]. These carbons are usually produced by sputtering from a solid carbon target. This is done in a process called “physical vapour deposition” (PVD). This is a trade-off between imparting enough energy to facilitate a transition from the low-energy, low-density graphitic state (sp^2) to the high-energy, high-density diamond state (sp^3), while not transferring enough energy to allow the material to undergo a spontaneous transition back to the thermodynamically preferred graphitic state.

Nonetheless, hard DLC films with very high sp^3 content (~80–90%) have been produced [75]. As mentioned earlier, the work by Merkulov *et al.* [71] focused on correlating the Raman spectra of diamond-like amorphous carbon thin films with varying sp^3 contents as measured by EELS. However, clear quantifiable conclusions regarding the Raman spectra and the sp^3 content was not possible. This may have been caused by clustering of the sp^3 -bonded atoms, potential errors in the sp^3 content measured by EELS or damage to the molecular structure during the analysis. The work by Knight and White [67] indicates that the stress state of the deposited diamond film also affects the observed Raman spectra. The stress in the deposited material is determined by the choice of

substrate. All of these factors make direct quantification of the sp^3 content of DLCs very difficult.

However, we can now place these structures on a ternary phase diagram and start to form an overall picture of amorphous carbon. To represent the spectrum of possible carbonaceous materials fully, it is necessary to add the third component, hydrogen, to complete the overview of amorphous materials. This allows the inclusion of hydrocarbon (HC) polymers, graphitic precursors and organic molecules, as can be seen in Figure 2-41.




Figure 2-41: Ternary phase diagram of amorphous carbons (after Ferrari & Robertson [76])

Thus a continuous transition from graphite to amorphous carbon is possible, whereas, due to the extreme hardness of diamond, such a gradual transition to amorphous material is not possible. It may be possible to use ion bombardment of diamond, but this would affect only very small regions, making bulk characterisation using Raman spectroscopy difficult. Instead, the approach to diamond is found in the region of increasing crystallinity of the PVD-grown DLC. As one moves from the soft amorphous carbon produced from graphitic material, the hardness increases as the amount of sp^3 or diamond-like material increases [77]. As the hydrogen content increases, the material again becomes soft as the structures progress towards more polymer-like materials.

It is clear that the Raman spectra of amorphous materials are not simply intermediate blends of the diamond and graphite spectra. The situation is further complicated because some peaks are enhanced over other peaks due to, for example, resonant Raman coupling [71];[76] which enhances the intrinsic intensity of the graphitic peaks. Visible Raman spectroscopy is therefore 50–230 times more sensitive to sp^2 sites, as visible photons preferentially excite their π states. Another factor that clouds the transition is the tendency towards clustering of specific bonds [76]. Raman investigation of these transition materials is a very complex field. A detailed analysis has been done by Ferrari and Robertson [76], the main points of which will be summarised here. A comprehensive model for the interaction of Raman scattering by amorphous carbons is derived based on an “amorphisation trajectory”. Hereby the progression of the Raman spectrum of amorphous carbon is described by the systematic transformation of graphite to tetrahedral amorphous carbon, via the intermediate stages of nanocrystalline (NC) graphite and sp^2 -hybridised amorphous carbon (*a*-C). This transformation is summarised as follows:

Stage 1: from graphite to NC-graphite

1. The G peak moves from 1581 to $\sim 1600\text{ cm}^{-1}$.
2. The D peak appears and $I(D)/I(G)$ increases following the relationship given by Tuinstra and Koenig in Section 2.3.5.3.
3. There is no dispersion of the G mode.

Stage 2: from NC-graphite to *a*-C

1. The G peak decreases from ~ 1600 to $\sim 1510\text{ cm}^{-1}$.
2. The relationship given by Tuinstra and Koenig is no longer valid.
3. Increasing dispersion of the G peak occurs.

Stage 3: from *a*-C to *ta*-C

1. The G peak increases from ~ 1510 to $\sim 1570\text{ cm}^{-1}$.
2. $I(D)/I(G)$ is very low or zero.
3. Dispersion of the G peak occurs.

Based on this analysis a few important conclusions can be drawn. The postulate proposed earlier is confirmed: during the initial stages of amorphisation there is an increase in the $I(D)/I(G)$ ratio, but this occurs without the presence of

any sp^3 -hybridised material. That is to say, the emergence of the D peak is a function of the breakdown of the sp^2 crystal structure, not of the creation of sp^3 bonds. The relationship proposed by Tuinstra and Koenig has a limited region of applicability and should be applied only in the region where pre-existing crystallites are being broken down.

Furthermore, when a reverse “ordering” procedure was attempted, it was found that the Raman spectra exhibit a type of hysteresis, for example during high-temperature deposition of *ta*-C or post-deposition annealing. During this transition two key processes occur, namely conversion of sp^3 bonds into sp^2 bonds, and clustering of the sp^2 material eventually to form ordered ring structures. The relative contributions of these processes may be skewed, depending on the exact thermal history of the sample. Thus due to the high sensitivity of the Raman spectrum to clustering via the resonance effect, there is no unique relationship between the $I(D)/I(G)$ ratio or the G peak position and the sp^3 content.

This has significant implications for the use of this approach to quantify domain sizes and molecular structure in synthetic graphite materials. These materials are derived from a gradual increase in the graphitic ordering of amorphous precursors. It is not possible to infer the relative sp^2 - sp^3 contents, and hence the corresponding molecular structures, since these depend on the different types of precursor used and their respective starting arrangements. Thus the Tuinstra and Koenig expression cannot be applied unless there is a large degree of certainty that the ordering process under consideration simply represents an increasing connectivity of discrete NC-graphite domains, as opposed to a transition from a mixed sp^2 - sp^3 molecular structure.

Furthermore, for a collection of carbon samples with an unknown history and starting structure, it is impossible to determine whether the variations in the $I(D)/I(G)$ ratio for the samples are caused by variations in the sp^2 crystal structural domains, or by the presence of an sp^2 - sp^3 mixed molecular structure within the material. This casts doubt on the validity of using the Raman spectra as a way of inferring the amorphous content of a sample.

Clearly, both the thermal and physical history of a sample are important in determining the amorphous structure of carbon samples from various sources. The carbon phase diagram is shown in Figure 2-42; it illustrates the fact that

graphite is the thermodynamically favoured phase at low pressure. Thus it is not unreasonable to expect that all amorphous carbon materials that have been heat treated for reasonable amounts of time at high temperatures (>1500 K) will exhibit a predominantly sp^2 -hybridised molecular structure.




Figure 2-42: Carbon phase diagram (from Burchell [77])

In the previous discussion the transition from amorphous material to graphite was considered as always being possible, but this is not always the case. At this point it is important to consider the differences between graphitising carbons and non-graphitising carbons. Certain carbon materials do not graphitise readily when heated to high temperature (>3 300 K) at normal pressure. These carbons form hard low-density materials, which can have very large surface areas due to their microstructure or even nanoporous structure. Typical examples are the so-called “carbon blacks”, activated carbon and glassy (or vitreous) carbons. Carbon blacks are soot or char materials which are generally used as a filler in rubber products or as pigments. Activated carbons are a special variety of carbon blacks which may undergo additional chemical activation through the use of steam or CO_2 to yield very high surface areas; they are used as gas and liquid adsorbents. Due to their chemical inertness, thermal shock resistance and good electrical conductivity, glassy carbons are used as reaction vessels and crucibles for chemical processing, as well as for acid-battery electrodes.

Suitable non-graphitising carbon materials are on the whole polymeric in nature, with the following characteristics:

- The polymeric backbone must be cross-linked.
- Carbonation must occur in the solid state, i.e. without mesophase formation.
- In general, molecular weight and degree of aromaticity have to be high to ensure a high carbon yield.

Typical precursors for these materials are polyfurfuryl alcohol, phenolics, polyimide, polyacrylonitrile (PAN), cellulose, polyvinylidene chloride, polyvinyl alcohol, polyphenylene oxide and aromatic epoxies. Thus PAN-based carbon fibres also fall into the general category of non-graphitising carbons.

From TEM studies of carbonised materials, the differences between graphitisable and non-graphitisable materials are readily observable. During low-temperature carbonisation (~1 000 °C) non-graphitisable materials form isotropic structures that are highly disordered, whereas graphitisable carbons show a clear tendency towards preferred orientation. When heated to graphitising temperatures, these materials form small graphitic domains, of the order of less than 10 nm in length.

The majority of these materials contain sp^2 -hybridised carbon [66]. Since the earliest work of Franklin [78] the tendency has been to think of the atomic structure of these non-graphitising materials as being composed of regions of multilayered graphite crystallites. Franklin's own model, shown in Figure 2-43, suggests small graphite crystallites connected via cross-linked branches. Several variations on this model have been suggested, such as the ribbon-like model of Ban *et al.* [79], shown in Figure 2-44, or the randomly stacked sheet-like model of Crawford and Johnson [80], shown in Figure 2-45.




Figure 2-43: Franklin's model of (a) non-graphitising and (b) graphitising carbon [78]



Figure 2-44: Ribbon-like model of Ban *et al.* [79]




Figure 2-45: Sheet-like model of Crawford and Johnson [80]

Based on recent evidence, Harris [66] has proposed a different molecular structure based on fullerene-like building blocks. This model proposes a far more disordered structure for non-graphitisable carbons. Using a Monte Carlo-type simulation, Kumar *et al.* [81] simulated the evolution of an amorphous polymer

into an amorphous microporous carbon structure. This model provides some insight into the evolution of the amorphous carbon structure as proposed by Harris. However, an undisputed picture of the molecular structure of amorphous carbons is still the subject of ongoing research.




Figure 2-46: Kumar *et al.*'s simulation of the structural evolution of amorphous carbon [81]

Thus the model proposed by Harris is capable of consisting solely of sp^2 -hybridised carbon and yet having a random, disordered structure. This further complicates the prior discussion since it implies that even for structures that have 100% sp^2 content, the layered, graphitic nature of the structure is not guaranteed. This adds further doubt to the application of the Tuinstra and Koenig expression (Eq. 2.5) for carbon materials that have undergone high-temperature graphitisation.

From this section and the previous section it should be clear that the molecular structure of imperfect carbon materials is an extremely complex subject. It depends on a very wide variety of factors and it is difficult to link the molecular structure directly to measurable characteristics. Hence it is difficult to come to absolute conclusions regarding the potential reactivity of these materials based on their nature, i.e. amorphous, diamond-like and graphite-like. This has very important implications for the reactivity of carbon materials as will become clear in the following sections, especially those when considering the active surface area of these materials.




2.5 The reaction of oxygen with carbon

The combustion of carbon and its hydrocarbon derivatives is arguably the most important reaction known to mankind. Since the earliest ancestors of man discovered the use of fire, combustion has provided the world's energy needs. Even today, despite large efforts to utilise renewable energy sources, 78% of the world's energy consumption is still provided by the burning of fossil fuels [82]. The combustion of coal alone accounts for almost 30% of this need [83]. Thus an exceedingly large volume of literature is dedicated to furthering the understanding of this reaction, complicated even more by the plethora of possible carbon allotropes that combust.

The reaction of carbon with oxygen is not the only reaction of industrial significance. The gasification of carbon with steam produces carbon monoxide and hydrogen, which can then be converted to virtually any hydrocarbon or organic chemical using, for example, the catalytic Fischer-Tropsch process. Activated carbons are produced by the activation of carbonaceous materials through the use of steam, carbon dioxide or air. Hydrogen can also be reacted with carbon to produce methane, although this reaction is of less industrial significance. All of these reactions are fairly similar and interrelated.

Despite the fact that gas-carbon reactions are an integral part of the industrial economy, fundamental understanding of their reaction mechanisms and kinetics has lagged far behind their practical use. This is due primarily to the lack of experimental techniques to define the properties of one of the reactants, namely carbon [84]. This has led to an exceedingly large amount of differing and apparently conflicting data for the order and other kinetic parameters of the carbon-oxygen reaction. It is interesting to note that Gulbransen and Andrew, two of the early investigators of graphite oxidation, in a paper written in 1952 [85] remarked: "The reaction of carbon with oxygen has been the subject of many studies over the past 100 years. A review and evaluation of this literature would be an extremely long and arduous task". In the subsequent 60 years it appears that the number of studies conducted has risen exponentially, yet a comprehensive review is still not forthcoming and is probably an impossible task at this point.

As a basic starting point, however, the four primary reactions of importance during the carbon-oxygen reaction are shown below.



The first three reactions are exothermic, while the last is endothermic. It should be clear that these four reactions are not independent as reaction (2.9) may be obtained by adding the reverse of reaction (2.6) to twice reaction (2.7). Furthermore, reaction (2.8) is obtained by adding reaction (2.6) to the reverse of reaction (2.7). From thermodynamics it can be shown that the equilibrium constant is defined by Eq. (2.10) [86].

$$\ln K = \frac{-\Delta G^o}{RT} \quad (2.10)$$

By taking the temperature derivative of this expression, the temperature dependence of the equilibrium constant is found to be:

$$\frac{d \ln K}{dT} = \frac{\Delta H^o}{RT^2} \quad (2.11)$$

If one assumes that the standard enthalpy change (heat) of reaction is independent of temperature, Eq. (2.11) can be integrated to give:

$$\ln \frac{K_1}{K_2} = -\frac{\Delta H^o}{R} \left(\frac{1}{T_1} - \frac{1}{T_2} \right) \quad (2.12)$$

Since K indicates the equilibrium composition, a plot of K vs. T will determine to what extent the forward or reverse reactions are favoured at a given temperature and pressure. Using Eq. (2.12) and Table 2-11, the plots of the equilibrium constant against temperature for the reactions given above are as shown in Figure 2-47.

Table 2-11: Standard Gibbs free energy and enthalpies of reaction
(taken from Smith *et al.* [86])

	ΔG° (kJ/mol)	ΔH° (kJ/mol)
Rxn (2.6)	-393.5	-394.4
Rxn (2.7)	-110.5	-137.2
Rxn (2.8)	-283.0	-257.2
Rxn (2.9)	172.5	120.0

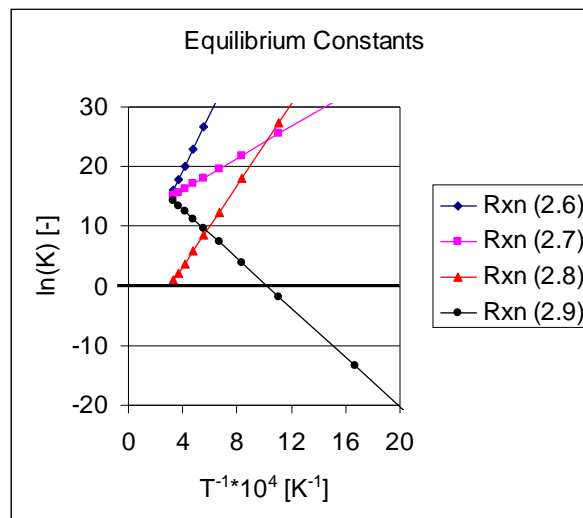


Figure 2-47: Temperature dependence of equilibrium constant

This figure illustrates a few noteworthy points. The equilibrium constants for reactions (2.6), (2.7) and (2.8) are always positive and always increasing. This is not unexpected for exothermic reactions and indicates that the forward reaction is increasingly favoured at lower temperatures. Reaction (2.9) exhibits an interesting behaviour in the sense that below 1 000 K, the equilibrium constant becomes less than one and the reverse reaction is favoured. Hence carbon monoxide decomposition to form carbon and carbon dioxide is preferred.

However, this reaction is extremely slow and attainment of conversion requires the presence of metallic catalysts [87].

At these temperatures the forward reactions of all three of the other reactions are large and positive, indicating that the forward reactions are heavily preferred. Thus under oxidative conditions, equilibrium considerations show that combustion is the preferred reaction pathway and one would expect high carbon dioxide concentrations and very little carbon monoxide.

A major issue regarding the combustion of carbon is whether the carbon dioxide is formed as a primary product or via the gas-phase oxidation of carbon monoxide. Walker *et al.* [84] examined the results of several researchers and concluded that there was near-total agreement that both carbon monoxide and carbon dioxide were primary products of reaction. However, this has been contradicted by recent work on fluidised beds [88];[89] which indicates that only carbon monoxide is the primary product. On the other hand, this work has been refuted by a recent study [90] in a similar set-up, but using different particle sizes and low oxygen concentrations.

However, the original study by Walker *et al.* [84] was not without uncertainty, specifically with regard to the exact value of the CO:CO₂ ratio, especially as the function of temperature remains contentious. They summarise their conclusions as follows:

- Both carbon dioxide and carbon monoxide are primary reaction products.
- The ratio of these products, CO:CO₂, generally increases steeply with increasing temperature [91].
- The ratio does not appear to be a function solely of temperature, i.e. it is dependent on the carbon reacted.

Perhaps this last point may explain the observed differences in the studies mentioned above since the former was conducted on a non-porous graphite while the latter used a bituminous coal char. This may be explained in terms of differences in surface heterogeneity, which will be covered in more detail in Section 2.7.

The reaction of carbon monoxide with carbon has been shown to occur at very low rates close to 700 °C [92]. This is in agreement with the earlier statement from thermodynamic considerations, namely that there may be a

reversal in the reaction direction around this temperature. Furthermore, experimental evidence indicates that the presence of carbon monoxide retards this reaction [93]. This has led to the development of several possible, complicated reaction mechanisms similar to the conventional Langmuir-Hinshelwood expression. However, there is significant disagreement regarding the exact reaction pathway [84].

The mechanisms given by Eqs (2.6) and (2.7) represent vastly simplified approximations of the real reaction pathway. In fact, it has been known for a long time that the reaction of carbon with oxygen consists of several intermediate steps, including the chemisorption of oxygen [85];[94]. The oxygen forms a surface complex with carbon, which cannot be removed. When carbon is degassed at high temperatures, only carbon monoxide or carbon dioxide is released. These complexes and their role in the reaction will be discussed in more detail in Section 2.8.

Under conditions where mass transfer limitations are completely alleviated and kinetics control the reaction, a qualitative understanding of the reaction order for a gas-solid reaction is possible. For simplicity, the reaction rate is assumed to be determined by the rate of desorption of the surface complex to form the gaseous product. The argument will be the same if the rate-limiting step is the rearrangement of surface groups to form a product that is readily desorbed. In this case, if the fraction of the surface covered by the complex is x , the reaction rate is proportional to x . At a particular temperature, if the surface coverage is nearly complete, i.e. $x \rightarrow 1$, the rate will be nearly independent of the oxygen pressure and the reaction order will be zero. On the other hand, if $x \rightarrow 0$, the change in surface coverage will be directly proportional to the change in pressure. At intermediate surface coverage, the reaction order will have some intermediate value. It should be clear that the converse is also true: as the pressure is increased, surface saturation will occur and $x \rightarrow 1$. The reverse is also true: as pressure is lowered, $x \rightarrow 0$. Thus the reaction order at a given temperature may range from zero to one, depending on the range over which the pressure is varied.

This is an oversimplification but provides a theoretical background for assessing the observed reaction orders. In general, the reaction order of carbon with oxygen has been found to be close to first order [84]. However, there are

several exceptions for which an explanation has been proposed by Hurt and Haynes [95]. Using surface heterogeneity considerations and a distribution of desorption energies, rather than discrete values, it may be shown that the reaction is nominally first order in oxygen. Furthermore, a wide range of fractional orders reported in various other investigations are explained on the basis of a calculated global reaction order. This observed order is based on a combination of first-order adsorption-limited reaction on bare sites and mixed-order reaction on partially covered sites.

The situation is further complicated by, for example, the existence of more than one type of surface complex [96]. For coal chars the presence and influence of the ash may also be responsible for some of the observed variation [97]. Finally, the influence of mass transfer limitations is sometimes difficult to avoid and some investigators incorporate the traditional effectiveness factor [98] in an effort to compensate for these effects. However, it is clear that unless the true kinetic pathway can be established, the reaction order of choice will most likely remain an empirical n -th order-based approximation [99]. These mechanisms will be considered in more detail in subsequent sections.

Studies have focused on a very wide variety of carbon materials and apparent activation energies with many variations have been reported [84];[100];[101]. The reasons for these variations may be anything from the influence of surface heterogeneity to mass transfer limitations, to the presence of impurities. However, Smith [102] showed that for an exceedingly wide variety of carbons, their intrinsic reactivity could be explained using a single activation energy of 179 kJ/mol, as shown in Figure 2-48.




Figure 2-48: Intrinsic reactivity of carbon as presented by Smith [102]

2.6 The role of mass transport and pore structure in gas-carbon reactions

In general, the heterogeneous reaction rate for a porous solid and a gas may be controlled by one of three major steps:

1. Mass transport of the reacting gas and gaseous reaction products across a relatively stagnant gas film between the exterior surface of the solid and the main gas stream
2. Mass transport of the reacting gas from the exterior surface of the particle through the pores to the reaction site beneath the surface; at the same time gaseous reaction products are transported in the opposite direction
3. Chemisorption of the reactant, wholly or in part, followed by rearrangement of the surface species to desorbable products and finally, desorption

Comprehending the role of these three steps is imperative for gaining an understanding of the kinetics of the gas-carbon reactions. Ideally, the functioning of the previous mechanism has three temperature-dependent zones [103]. At low temperature, Zone I, the reaction rate is governed solely by the reactivity of the solid. The rate is low enough for all species to have enough time to diffuse across

any barriers and into the pores, resulting in an even concentration of reactant gas (and products) everywhere. The activation energy that is measured is the true activation energy. The effectiveness factor, η , is defined as the ratio of the measured reaction rate to the true reaction rate if the gas concentration, C_g , is constant (or rather has an insignificant concentration gradient), i.e. the entire surface area is exposed to the bulk gas concentration. In Zone I the effectiveness factor is one. A schematic representation of the reaction zones and gas concentration is shown in Figure 2-49.

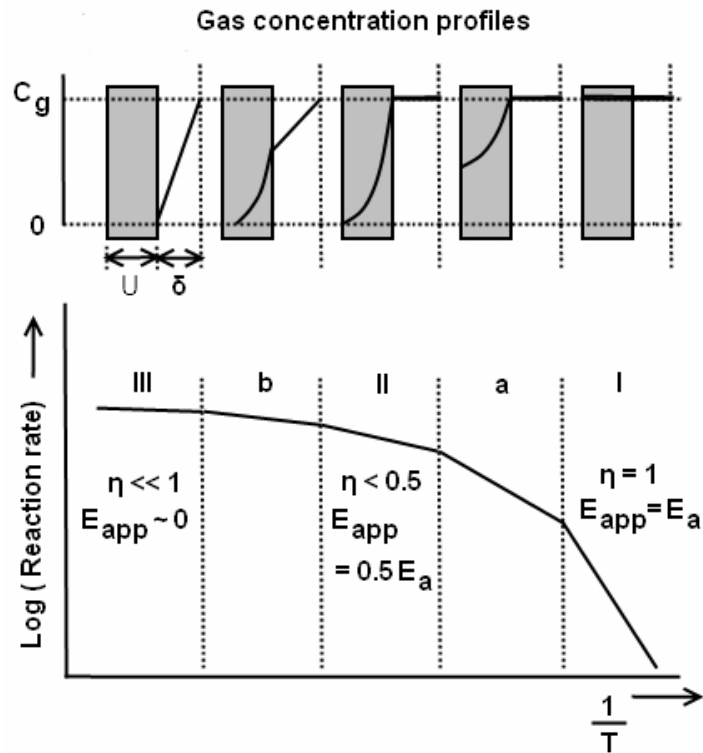


Figure 2-49: Gas-carbon reaction zones

In the intermediate region, Zone II, the reaction rate increases and the gas concentration falls to zero within some distance less than the radius of the particle, R . It can be shown that in this zone the apparent activation energy reduces to roughly half of the true activation energy [103] and the effectiveness factor becomes less than one half. In the high-temperature range, Zone III, the reaction rate is now so high that the reactant gas does not penetrate the particle and the reaction proceeds almost exclusively at the outside surface. The reactant

gas concentration falls to some small value through the stagnant boundary layer (δ) that surrounds the particle. Hence in this zone the reaction rate is controlled by Step 1. Increasing the temperature only affects the reaction rate insofar as it allows more reactant gas to reach the surface per unit of time. Since bulk mass transport has low activation energies, the apparent activation energy is also low and the effectiveness factor is obviously far less than one.

There are several reasons why this simplified representation of the temperature behaviour of the reaction rate may not be accurate in practice:

1. The concentration gradient across the stagnant boundary layer may deviate significantly from zero before the reactant concentration reaches zero somewhere within the solid. In this case Zone II will effectively disappear. This may occur under conditions where the bulk gas flow rate is low and the particle size is small, so where the ratio of external surface to volume is high.
2. The rate-controlling part of Step 3 in the mechanism may change as a function of temperature. For example, if the rate-controlling step changes from adsorption in Zone I to desorption in Zone II, the true activation energy for the reaction will have changed. However, the apparent activation energy measured in Zone II will correspond to only half its true value, which will no longer correspond to half of the activation energy measured in Zone I.
3. In Zone II the concentration of product gases may become significant, leading to significant retardation of the reaction. In this case the true activation energy becomes a complex mixture of activation energies for the different reactions. In addition, the complicated nature of multicomponent gas diffusion [105] makes this case difficult to model accurately.
4. Most porous solids do not have a single, narrow pore distribution. This is especially true of carbon materials as these may have varying combinations of macro-, meso- and micropores [106];[107]. Thus it is possible to develop a variety of concentration gradients within the porous solid, depending on the size distribution and interconnectivity of the pore network.

These factors can make it difficult to determine the true activation energy of a given sample without *a priori* knowledge of where these transitions occur. This, coupled with the complexity of the gas-carbon reaction, i.e. multiple reactions for multiple reaction products, has led to the development of complicated two-film models [108];[109] to explain the observed behaviours. Furthermore, models similar to the traditional shrinking core model have also been suggested [110];[111] to account for ash accumulation. Recent advances in computational flow dynamics have allowed the complex modelling of the flow and diffusion regions surrounding a carbon sample [112];[113]. However, each modelling effort utilises a different kinetic pathway. If mass transfer considerations are to be taken into account, it is clear that, as before, the true kinetic pathway must first be established since without this, the concentration profiles developed under mass transfer-limited conditions remain speculative at best.

The reaction rate will also depend heavily on the exposed surface area or pore structure of the carbon. It has long been known that the microtexture of coal influences its reactivity [114]. However, as mentioned earlier, the pore structure of carbonaceous materials can be exceedingly complex. Most often gas adsorption and hence adsorption isotherms are used to characterise carbon pores. However, these methods depend on complex models which often lead to erroneous pore diameters when applied to carbons [115]. Surface heterogeneity also presents significant challenges to developing accurate models for these pores [116].

Furthermore, carbon materials can have a wide variety of different pore structures, which further complicates the interpretation of the isotherm data. For activated carbons these may be broadly classified as [117]:

- *Molecular sieve-type pores*. These are narrow micropores (<0.7 nm) and have a sharp pore size distribution.
- *Representative pores*. This type is representative of activated carbons. They have a sharp, major, small micropore distribution and an additional wide, minor, larger micropore distribution.
- *Bimodal pore distribution*. These pores have two clear, fairly even regions of larger and smaller micropores.

- *Super active carbons*. These carbons undergo intense activation to produce highly porous materials with a super-high surface area, with pore widths around 2 nm and a medium to wide distribution

These tests give little information regarding the interconnectivity of these pores and few indications of restricted access pores. Modelling the structure and evolution of these pore networks during gasification poses a significant challenge. A model that has demonstrated significant popularity is the *random pore model* [118]. It was originally successfully applied in the 1980s to model the oxidation kinetics of several different coal chars [119]. Most significantly, it has provided insight into the manner in which the pore size distribution grows as a function of conversion. It has successfully predicted the experimentally measured pore size as a function of conversion for coke and anthracite char [120].

Since then the model has been applied numerous times [121] and adapted in several ways. It has been discretised [122] to account for possible errors in its application to highly microporous materials. In this case the pore sizes are only a few multiples of the spacing between the reacting sites in the char. Thus the concept of a continuously moving interface is inappropriate. This version of the model allows for the incorporation of regions with different initial surface chemistries and different reactivities, based on the orientation of the underlying crystallites [123]. This approach has been successfully applied to coal chars [124].

Real particles undergoing combustion under practical conditions generally undergo additional effects. For example, under certain conditions particle disintegration may occur to a significant extent. Modifications to the random pore model have been proposed to account for this phenomenon during microporous char combustion [125]. The presence of metallic catalysts, especially alkali and alkali earth metals (AAEMs), can significantly affect the oxidation behaviour of carbon materials. Several modifications to account for the action of these impurities have been proposed and applied to coal char [126] and a very wide variety of biomass [127].

These effects illustrate the complexity involved in modelling the oxidation of carbon materials where mass transfer limitations may be present or the development of the microporous surface area dictates reactivity. A final

consideration that should be borne in mind is the relative crystalline order of the carbon sample under investigation. It has been shown that even under heat treatment to relatively low temperatures ($<1\ 400\ \text{K}$) samples undergo significant reductions in the pore size and changes in the structure due to crystallite ordering [128]. In addition, this effect is extensively influenced by the presence of metallic impurities [26]. This structural ordering of the crystallites in the pore wall will have a direct influence on the reactivity of the sample, since the basal planes of these crystallites are comparatively inert. At the same time, the overall surface area may remain largely the same. This leads to the notion of the “active surface area”.

2.7 The concept of active sites and active surface area

As mentioned in the previous section, the random pore model has been utilised to model the oxidation of a variety of carbonaceous materials, mostly chars. It models the growth and coalescence of pores within the material and hence is a representation of how the surface area of the char particle develops. This implies that the reaction rate is heavily dependent on the surface area. In general, this is found to be true for amorphous carbons, but as the degree of graphitisation, or rather the crystalline order, of a carbon increases, this assertion becomes less accurate.

The cause lies in the atomic structure of the graphite crystal. From a schematic representation of the graphite crystal, shown again in Figure 2-50, the reason for the reactive anisotropy becomes clear. The crystals' basal atoms are linked in three directions with adjacent carbon atoms. These form very strong covalent σ bonds with each other and hence are comparatively difficult to remove from the lattice. This is further accentuated by the fact that there are additional π bonds between these atoms and thus the carbons' electronic requirements are fulfilled. These atoms share four electrons with other atoms, filling their outer shell completely. Hence the basal plane of graphite may be considered inert up to very high temperatures and it is exceedingly difficult for gaseous reactants to bond to these atoms and form a desorbable product.




Figure 2-50: Reaction anisotropy of the graphite crystal structure

It should be borne in mind that this is the basal plane of the ideal, perfect graphite crystal. Crystal and lattice defects will dramatically influence the reactivity of the basal atoms around it [129]. Furthermore, the planar, layered structure of graphite shields atoms located in successive basal planes below the outer layer, making it very difficult for reactive gases to gain access to these atoms to attack and remove them.

The edge atoms, however, are a completely different matter. They are exposed at all layers in the structure to attack from gaseous reactants. In addition, these atoms are less tightly bound to the crystal structure since there are only two covalent σ bonds to the adjacent carbon atoms. The electronic requirements of these atoms are not fulfilled and they can be easily bound by gaseous reactants.

The atomic structure shown in Figure 2-50 is potentially incomplete as it represents a dangling carbon atom. These dangling atoms have an unpaired electron, potentially two depending on the nature of the π bond structure of the adjacent atoms. Such an atomic configuration is potentially stable under inert conditions [60], but would be highly reactive in the presence of oxygen. Thus, in general, these atoms are presumed to be bound to a variety of oxygen-based surface complexes. The exact nature of these complexes is not currently known and this subject is discussed in more detail in the next section, Section 2.8.

From Figure 2-50 it is also clear why the surface area of the particle is no longer representative of its reactivity. The ratio of edge atoms to basal atoms is a strong function of the crystallite size. The representation of amorphous carbons envisioned by Franklin [78] (Figure 2-43) provides some visualisation of the transition from amorphous to graphitic carbons. As the crystallites grow in height

and width from a few Angstroms to a few microns, the ratio of edge to basal atoms will decrease dramatically. Correspondingly, the total surface available for reaction will become smaller and it will also become a progressively smaller fraction of the total surface area. The overall surface area will also decrease as pointed out in the previous section by the shrinkage of the pore walls [128] as a function of heat treatment.

This concurrent reduction in the total surface area implies that for most carbon materials there is still a rough relationship between surface area and reactivity, as pointed out by the work of Smith [102] and illustrated by Figure 2-48. In this figure the reaction rates for the different carbon materials are normalised based on their surface area.

However, as pointed out by Smith, there is still a large difference in the reactivities of various carbon materials despite the removal of surface area effects. It is postulated that these differences may be caused by differences in the atomic structures and the impurity levels of the carbons. This is reinforced by other studies [130];[131] which show a lack of correlation between surface area and reactivity for some carbons.

The concept of “active surface area” was originally developed by the group of Phillip Walker Jr. at Pennsylvania State University in 1959. Early speculation indicated that the active surface area of a sample was not only a function of crystallinity, but also a function of crystallite orientation, lattice defects and impurities [84]. Hence simply using XRD to determine the degree of crystallinity is not sufficient and a new method for measuring the active surface area had to be developed. Conventional surface area measurement utilises inert gases such as argon and nitrogen, which are also not suitable for discerning between chemisorption (active sites) and physisorption (BET surface area).

The technique of using thermal desorption was first applied successfully by Laine *et al.* in 1963 [132] to measure the development of the active surface area as a function of conversion for Graphon, which is a highly graphitised carbon black, Sphero carb. A sample of Graphon was oxidised to seven different levels of conversion between zero and 35% at a temperature of 625 °C. The samples were then heated to 950 °C under vacuum for 3 h to ensure that any surface complexes present from the original oxidation were desorbed. Prior

studies [133] had shown that 950 °C was sufficient to fully desorb the surface complexes present.

The total surface area (TSA) was measured using the BET method and nitrogen adsorption. The active surface area (ASA) of the sample was determined by first exposing the heat-treated sample to oxygen at 300 °C for 24 h at oxygen pressures of around 70 Pa. Then the oxygen was evacuated, the sample heated to 950 °C and the amount of oxygen recovered measured using an MS. In order to calculate the ASA, the assumptions were made that each complex consists of one oxygen atom per carbon edge atom and that each carbon atom occupies an area of 0.083 nm².

The samples were then oxidised once more, in a closed vessel, at temperatures of 575, 625 and 675 °C and oxygen pressures of around 5 Pa, so-called “reaction rate runs”. The amount of additional burn-off during these runs was kept very low (< 0.1%) to ensure that the TSA and ASA remained constant. During the run the CO and CO₂ formation was continuously monitored and a mass balance was used to calculate the fraction of ASA open for reaction and the fraction occupied by a stable oxygen complex. Thereafter the samples were again heated to 950 °C under vacuum to recover the oxygen complex on the surface and confirm the accuracy of the mass balance.

The basis of this investigation is the simplistic kinetic pathway given by Eqs (2.6) and (2.7). The reactivity of the carbon is then determined solely by the fraction of active, or edge, sites that are unoccupied by a stable surface complex formed between carbon and oxygen, although the reaction is thought to occur via a fleeting complex. First-order dependence on oxygen pressure is assumed and the reaction rate expression is given as:

$$-\frac{dP_{O_2}}{dt} = k'_{O_2} P_{O_2} TSA = k_{O_2} P_{O_2} ASA(1 - \Omega) \quad (2.13)$$

where P_{O_2} is the oxygen partial pressure, Ω is the fraction of the ASA that is occupied by the stable surface complex and k'_{O_2} and k_{O_2} are the reaction rate constants based on TSA and ASA respectively.

This study revealed several important points:

- During the reaction rate runs, it was found that the fraction of the ASA occupied by the stable surface complex gradually increased until saturation coverage was reached, with the time until saturation is achieved decreasing with increasing temperature.
- The reaction rate constants based on the TSA show a sharp initial decrease before levelling off, whereas the rate constants based on the ASA changed relatively little during the reaction.
- The rate constants based on the TSA, calculated at the end of the reactivity run, showed a sharp monotonic increase with prior burn-off. In contrast, the rate constants based on the ASA were essentially constant.
- The amount of surface complex formed at saturation coverage increased with decreasing temperature, but this also increased the time required to achieve saturation coverage. This led to the choice of 300 °C as the point at which saturation coverage could be obtained over a practical time period, i.e. 24 h, for the initial ASA determination.
- The amount of surface complex formed at saturation coverage also showed a marked increase with prior burn-off, but the fraction of ASA occupied by this complex decreased with an increase in prior burn-off.
- During the desorption step of the initial ASA determination, almost all of the surface complex was recovered as CO and only a minor component as CO₂.
- The TSA of Graphon increased less than two-fold upon oxidation to 35% conversion, whereas the ASA increased by a factor of around 18, thus indicating the vast discrepancy between ASA and TSA development in highly graphitic samples.
- Following a reaction rate run, virtually none of the surface complex remaining on the surface could be desorbed by degassing under vacuum at the reaction temperature.
- At the highest temperature (675 °C) the reaction rate constants based on TSA are not initially fully normalised and time independent.

These initial observations already hint at the complexity that underlies the surface reaction of carbon with oxygen. The existence of both stable and unstable surface complexes significantly complicates the kinetic mechanism. Their characteristics are made even more intricate by the fact that they show both temperature and conversion dependencies, the latter possibly indicating surface heterogeneity effects.

The results of this study have since been confirmed by numerous investigations [134];[135] for a variety of carbons, including Saran char (a copolymer of vinylidene chloride and vinyl chloride), graphitised carbon black, spectroscopically pure carbon black and pyrolytic graphite. This asserts the validity of active surface area as a key indicator for carbon reactivity. However, some discrepancy has been observed for some bituminous chars [136], which indicates that chemisorbed oxygen is not always present in the form of the stable oxygen; this inhibits oxidation, even at low temperatures (400 °C).

The effect of heat-treatment temperature and the corresponding reduction in reactivity for an activated carbon is discussed and illustrated by Arenillas *et al.* [137]. They measured BET surface area, TPD-based ASA and XRD-derived crystallographic parameters to illustrate the evolution of structural ordering within the material. They conclude that BET surface area is not a relevant reactivity normalisation parameter and that ASA appears to be more fundamentally correct.

The last point raised by the study of Laine *et al.* indicates that at higher temperature some variability is still unaccounted for by the ASA and stable surface area concept. It was concluded that this was because at this temperature the oxygen complex was no longer truly stable and was perhaps acting as an intermediate for product formation. The stability and reactivity of these complexes were investigated by Vastola *et al.* [138] using the oxygen 18 isotope as a tracer. This study concluded that the nature of the stable surface complex was not dependent on the time of its formation, i.e. stronger adsorption taking place first, followed by progressively weaker sites.

Furthermore, the study showed that the stable surface complex formed during a given reaction run, which did not reach saturation, would be released during a subsequent reaction run, at the same temperature, using a different oxygen isotope. During this period the isotopically different stable surface complex would build up to saturation coverage. This result is contradictory to the

expected understanding that, once formed, the stable surface oxide will remain on the surface to block future reaction via that site. On the contrary, the result indicates a continual renewal of the “stable” surface complex. This work was continued by Walker and co-workers [139] and it was concluded that carbon dioxide formation proceeds via an intermediate which contains reversibly adsorbed oxygen.

Lussow *et al.* [140] continued to study the same Graphon studied by Laine *et al.* [132] and found that the assertion that the surface was completely covered by complex after 24 h of exposure to oxygen at 300 °C was false. The amount of surface complex present on the surface was dependent on both the exposure temperature and the pressure. Further work indicated that dissociative adsorption of the oxygen occurred onto two adjacent sites where the two atoms each form mobile surface oxides. However, to account for the additional adsorptive capacity it was necessary to introduce the concept of two types of active site with differing reactivities toward oxygen. This concept of multiple types of active site was supported in subsequent work by Bansal *et al.* [141].

The thermal decomposition of surface oxides formed during the low-pressure oxidation of Graphon was studied by Phillips *et al.* [142]. Their work asserted the veracity of the earlier statement that the stability of the formed surface oxide is dependent on its formation temperature. Furthermore, the desorption of carbon monoxide supports the idea of a limited number of discrete types of active site and that the carbon dioxide formed during desorption reacts readily with the carbon to form carbon monoxide. It is concluded that decomposition of the surface oxide results in the formation of highly reactive “nascent” sites with which the carbon dioxide reacts readily. This conclusion is further substantiated by Hall and Calo [143] who show that both carbon monoxide and carbon dioxide can undergo secondary interactions with either the oxygen complexes present on the carbon surface or free active sites. They assert that caution should be exercised when interpreting TPD spectra since these may not be representative of the original surface oxides, depending on the experimental conditions.

The presence of large amounts of stable surface oxides on carbon materials has significant implications for the use of only TGA measurements for determining the reactivity. As pointed out by several studies [144];[145], the

adsorption of surface oxides can contribute to large errors in the reactivity determined solely by weight measurement. Conditions for TGA-based kinetic studies should be carefully selected to ensure that chemisorption dynamics achieve a pseudo-steady state during the experiment.

Brown *et al.* [146] combined the use of an enclosed gas system connected to a mass spectrometer sampling system with a microbalance. Using this set-up they successfully monitored the surface complex development during reaction and subsequent TPD of the surface complexes formed during the reaction of three carbons with widely differing surface areas. In a later investigation [147], they appear to be the first group to have quantitatively measured and confirmed the existence of active sites with a continuous distribution of desorption activation energies, as opposed to the idea previously considered of a few discrete types of active site with different activation energies. This investigation also indicated that the distribution of activation energy is dependent on the temperature at which the oxygen was chemisorbed.

Following these studies, a fairly comprehensive review of the accumulated evidence was compiled by Walker *et al.* [148]. They evaluated studies on Graphon, Saran char, an alternative graphitised carbon black (V3G) and finely ground natural graphite, and reached several key conclusions:

- Fundamental, specific reaction rates can be obtained using high-purity, highly crystalline graphites.
- Obtaining fundamental rates for carbons with small crystallites and/or small particle size is significantly more challenging.
- The experiments done by Laine *et al.* [132] were conducted at very low oxygen pressure, which leads to negligible formation of the fleeting oxygen complex and only a small fraction of the ASA covered by stable complex.
- Studies conducted at much higher pressures show that the sites that chemisorb oxygen during gasification can significantly exceed the number of sites that chemisorb in the absence of gasification at low temperature. Thus the procedural approach of Laine *et al.* [132] does not have general applicability.

- Based on the possible explanations for the previously mentioned fact, it is concluded that oxygen chemisorption cannot be relied on to measure the total ASA of carbons.
- Strong site physisorption is suggested as an alternative, but the results are not conclusive.
- The assertion that the stable surface complex inhibits gasification is not necessarily true. Rather, the build up of this complex may be considered a necessary step to reduce the activation energy required for desorption of the fleeting oxygen complex. This is due to the coverage dependence of desorption energies.

In an effort to overcome these obstacles, Lizzio *et al.* [149] introduced the concept of “reactive surface area” (RSA). The RSA is defined as the fraction of the ASA covered in a carbon-oxygen surface intermediate that readily decomposes to form gaseous reaction products. It is noted that for some carbons, under certain conditions, either the ratio of RSA to TSA or that of RSA to ASA may be assumed to be constant. In order to measure the RSA, the usual TPD is utilised in conjunction with a new approach termed “transient kinetic studies” (TKS). A sample is partially gasified in carbon dioxide (flowing at 200 ml/min) up to a given conversion. Then the purge gas flow is switched to nitrogen (flowing at 200 ml/min) and the transient decay of carbon monoxide is monitored by a non-dispersive infrared sensor, at reaction temperature.

The TPD was carried out in a separate fixed-bed reactor coupled with an MS. The sample was gasified using carbon dioxide. Once it had reached the desired conversion, it was cooled. After the reactor system had been flushed with argon, the TPD run was done at 5 K/min up to 1 000 °C. This process was used to calculate the total amount of surface complex present. In order to determine the amount of stable surface complex formed at reaction temperature, the sample was gasified to the desired conversion and then the reactor system was purged with argon gas. Once carbon monoxide evolution had ceased, the sample was cooled and the same TPD run mentioned previously was done. It was not possible to use the MS to determine the exact amount of carbon monoxide evolved immediately after the change to inert gas. The second TPD spectrum is

then subtracted from the first to determine the behaviour of the fleeting surface complex.

The results for TKS and TPD are shown to be equivalent for three different chars (bituminous coal, Saran and lignite). Of slightly more relevance is the work done by Radovic *et al.* [150] on char gasification in oxygen. Unfortunately, they found it impossible to perform TKS on the bituminous coal under investigation. Under the gasification conditions studied, the decay of carbon monoxide evolution is extremely rapid and of the same order as the gas change to inert conditions. Thus it cannot be used to infer desorption rate kinetics. It is, however, suggested that under different conditions, i.e. higher temperature and lower oxygen pressure, it might be possible to achieve useable results. It is concluded that this low RSA is due to substantial coverage of the ASA by stable surface complex at the low reaction temperatures investigated.

This work fails to incorporate an analysis of either the fleeting or stable surface complexes that may be present at different temperatures, nor any possible transitions in the behaviour (i.e. stable to fleeting) that might occur. From previous evidence it seems unlikely that the behaviour of these species will be independent of temperature. Thus clearer knowledge of the chemical nature of these complexes is needed, specifically their atomic structure and stability behaviour as a function of temperature.

2.8 Surface complexes and possible kinetic mechanisms

Early *in situ* studies using FTIR spectroscopy indicated that carbon monoxide is formed from carbonyl surface groups and carbon dioxide from lactone groups [151] via the reaction of oxygen with carbon films derived from carbon suboxide. The appearance of both groups on the carbon surface initially occurs at room temperature. In addition, it was found that both carbon monoxide and carbon dioxide could be formed from acid anhydride surface groups [152] via TPD from Saran char oxidised in air at 375 °C and 1 atm pressure. A collection of possible surface groups is shown schematically in Figure 2-51.




Figure 2-51: Possible carbon surface groups

An extensive study on the desorption products of three graphite samples exposed to oxygen, carbon monoxide and carbon dioxide was done by Marchon *et al.* [153];[154]. The investigation utilised an ultra-high vacuum (UHV) chamber and an MS which monitored desorption products as a function of their atomic mass number (AMU). Some XPS spectra of the adsorbed species were also obtained. A combination of normal gases, oxygen 18 isotopes and carbon 13 isotopes was utilised. Gasification rates allowed kinetic parameters such as activation energy to be determined. The results from this study led to several useful conclusions:

- The shape, intensity and temperature dependence of the TPD spectra vary according to the temperature and pressure used for the initial adsorption; this is true for all three gases tested.
- Furthermore, the sticking probabilities are not constant with varying surface coverage.
- Carbon monoxide shows a very low sticking probability on a cleaned (1 500 K, under UHV) graphite surface.
- For exposure at low temperatures, carbon monoxide shows TPD peaks at 393, 503 and 673 K for AMU 28, whereas at high temperatures two peaks at 973 and 1093 K are observed.

- For exposure at low temperatures, carbon monoxide shows TPD peaks at 443 and 673 K for AMU 44, whereas at high temperatures two peaks at 673 and 923 K are observed. However the amount of carbon dioxide desorbed is always less than 10% of the total amount of carbon monoxide chemisorbed.
- For exposure at low temperatures, carbon dioxide shows TPD peaks at 923, 1 093 and 1 253 K for AMU 28, whereas for AMU 44 a single low-temperature peak is observed at 423 K. Carbon monoxide is the main desorption product.
- For exposure at low temperatures, oxygen shows TPD peaks at 463 and 693 K for AMU 44 and three peaks are observed for AMU 28 at 973, 1 093 and 1 253 K. At high temperatures three peaks at 573, 793 and 923 K are observed for AMU 44. The prominent desorption product is always carbon monoxide.
- Isotopic results indicate that carbon dioxide reacts with a bare active site to yield a surface complex (semiquinone) and carbon monoxide.
- Based on the activation energies of gasification and desorption, it is concluded that the reactions of carbon monoxide and carbon dioxide with carbon are desorption controlled.
- Low-temperature carbon monoxide adsorption yields weakly bound carbonyl and/or cyclic ether groups.
- High-temperature carbon monoxide adsorption and low-temperature carbon dioxide adsorption yield semiquinone or ketone groups. To desorb this species, which yields carbon monoxide, requires the breaking of two carbon-carbon bonds, which explains the high desorption temperature.
- Carbon dioxide is formed from desorption of a lactone group. This lactone group is possibly formed via a multistep reaction whereby carbon monoxide combines with an existing semiquinone or directly via carbon dioxide adsorption.
- The multiple TPD peaks for desorption of specific surface species are postulated to be due to active site heterogeneity, possibly due to differences between the armchair and zig-zag sites.

Brown and Haynes [155] studied the interaction of carbon monoxide and oxygen with Sphero carb. Their investigation reinforces several findings of the previous investigation, namely that carbon monoxide forms both stable and unstable surface complexes but the amount chemisorbed is very small compared with oxygen. The desorption of the stable complexes occurs only at high temperature and only a minor portion of the chemisorbed carbon monoxide is desorbed as carbon dioxide. One additional point of importance is the fact that the desorption of surface oxides created by oxygen adsorption is effectively irreversible.

Sibraa *et al.* [156] also found that only minor amounts of carbon dioxide formed during post-oxidation TPD and that the secondary reaction between carbon monoxide and the formed surface complexes is negligible.

The previous discussion aims to highlight the complexities involved in characterising the carbon surface structures and their behaviour. This point is further illustrated by Henschke *et al.* [157] who carried out an extensive battery of tests only to conclude that no meaningful kinetic law can be derived.

Zhuang *et al.* [158] combined TPD, TK, DRIFT and XPS spectroscopy to study the surface complexes formed during the gasification of a phenol-formaldehyde resin char, with better success. Their work indicated that during both TK and TPD carbon monoxide was the major product and that the amount of gas evolved during TK is much smaller than that released during TPD. This is in accordance with the work by Radovic *et al.* [150] discussed in the previous section which indicated that during reaction in oxygen, the fraction of ASA occupied by the stable surface complex significantly exceeds that covered by the unstable complex.

Carbon monoxide desorption peaks are in the 790–1 220 K range, while carbon dioxide desorption occurred at 930 K for a sample undergoing TPD immediately following TK. In the case where a sample is instead cooled from reaction temperature, in oxygen atmosphere, and then subjected to TPD, carbon dioxide release starts at lower temperatures, while carbon monoxide desorption occurs in roughly the same temperature range and is again much larger (~10 times) than that of carbon dioxide. However, later work [159] shows that some of the low-temperature carbon dioxide desorption is due to exposure to moisture,

which is presumed to transform lactone groups into less stable carboxylic acid groups.

From the DRIFT data it is concluded that the carbon dioxide-forming groups are lactones and/or acid anhydrides. Furthermore, carbon monoxide-forming groups are identified as strongly conjugated carbonyls (semiquinones) and/or ether, although the ether component may be due to the lactones or anhydrides. This work is in direct agreement with the study by Marchon *et al.* [154].

Subsequent gasification work by this group [160] at 773 K with isotopically labelled oxygen led to the proposal of the following kinetic scheme:

- The formation of carbonyl/ether complexes occurs with concurrent carbon monoxide production. Carbonyl/ether formation dominates lactone/acid anhydride formation.
- A major kinetic pathway is the secondary reaction of oxygen with a carbonyl/ether group to produce carbon dioxide and either carbon monoxide or the original carbonyl/ether group.
- A secondary reaction of oxygen with carbonyl/ether groups to form lactone/acid anhydride groups is known to occur, but quantification is not possible.
- The direct decay of lactone/acid anhydride groups to carbon dioxide and that of carbonyl/ether groups to carbon monoxide are not thought to be major contributors to the reaction rate.
- Both the secondary reaction of oxygen with existing surface groups and the direct decay of surface groups contributed to carbon monoxide and carbon dioxide released during isotopic TK studies. However, the dynamics for these two pathways appear to be different, leading to a fast and slow decay component in the TK spectra. This conclusion was reached during a switch in isotopic oxygen.

Based on this mechanism, Figueiredo *et al.* [161] developed a technique for the deconvolution of the TPD spectra. The primary decomposition products are assigned to the reactions stipulated by Zhuang *et al.* [158]. Using the carbon monoxide and carbon dioxide TPD spectra, a Gaussian multipeak function is

fitted to each spectrum. It is assumed that a normal distribution of desorption activation energies exists due to the surface site heterogeneity. Using the fit, the relative amounts of each surface complex can be calculated. This approach has been applied to a commercial activated carbon [162] and carbon nanofibres [163] to determine the effect of various treatments and functionalisations on the physical and catalytic properties of these materials.

A radically different type of surface complex was proposed by Pan and co-workers [164] based on investigations using a natural and nuclear graphite. Their work indicated that a large amount of adsorbed oxygen is released above the 1 000 °C maximum temperature used in most desorption studies. Samples were oxidised to 20% conversion in air at 600 °C. The BET surface area was determined using nitrogen adsorption and ASA was determined by TPD and using the assumptions of Laine *et al.* [132], i.e. that each complex consists of one oxygen atom per carbon edge atom and that each carbon atom occupies an area of 0.083 nm².

The experiments showed that for the nuclear graphite the ASA amounted to five times the BET surface area. Based on this, it was concluded that intercalation of dissociated atomic oxygen had occurred. Using molecular orbital theory, the most likely structure of the intercalated oxide was proposed to be as shown in Figure 2-52. This structure was chosen from four alternatives based on the highest calculated bond strength.




Figure 2-52: Intercalated oxygen surface complex

Although the ASA measured for the natural graphite was high relative to the BET surface area, given the assumption that these are perfectly crystalline flakes, it did not exceed the BET surface area. Furthermore, XRD results indicated no measurable change in the interlayer spacing and the Raman spectra showed only a small decrease in the I(D)/I(G) ratio. It is possible that oxidation of the nuclear graphite caused the exposure of significant amounts of edge planes

that were previously inaccessible [146], thus substantially increasing the fraction of the BET surface area available for chemisorption. This fact, in conjunction with errors in MS calibration and in the estimated surface area of the oxide, may have led to the above-mentioned observation.

Nonetheless, at around the same time Kapteijn *et al.* [165] were using TK and TPD to elucidate the mechanism for the reaction of carbon dioxide with activated carbon. They found, in similar fashion to Zhuang *et al.* [160], that isotopic TK studies (in this study they were conducted in an inert environment) resulted in a composite decay curve. The slow portion of the decay was initially attributed to a semiquinone species, while the fast portion was attributed to quick, reversible adsorption of a carbonyl species.

These two groups combined their work to publish a paper on the path to a “unified mechanism for carbon gasification”. Inspired by the structure for the surface complex suggested by Pan *et al.* [164], it was postulated that there are two main groups of active oxygen-containing structures on the edge sites: one is an in-plane group, represented by either semiquinone or carbonyl structures, and the other group contains both in- and off-plane oxygen atoms, in this case represented by singly bound oxygen atoms and in one case linked to a carbon atom that already has four bonds.

It was calculated, using semi-empirical molecular orbital theory, that the carbon-carbon bonds for the latter group are severely weakened and that they decompose at lower temperatures. The former group has stronger bonds which decompose only at high temperatures, above 950 °C. Hence the group containing oxygen bonded both in- and off-plane is responsible for the gasification reaction of carbon with oxygen at temperatures below 950 °C. It is also concluded that carbon monoxide desorption shows a shoulder below 950 °C only for the carbon oxygen reaction, which is in contradiction to the results of Marchon *et al.* [154]. Furthermore, Kapteijn *et al.* [166] directly call into question the assertion by Marchon *et al.* that carbonyl or ether surface groups are less stable than semiquinone or ketone groups. The basis for this assertion is molecular orbital calculations.

Moulijn and Kapteijn propose a fully unified mechanism [168] to explain the gasification of carbon by any oxygen-containing molecules, as shown schematically in Figure 2-53, in which C(O) represents a semiquinone group and

$C(O)C_f(O)$ represents the combined in- and off-plane surface complex, shown in Figure 2-54.




Figure 2-53: Unified carbon gasification mechanism (after Moulijn and Kapteijn [168])




Figure 2-54: Off-plane oxygen surface complex

Thus the stable oxygen surface complex plays a crucial role in gasification since it generates the site where the more reactive complex can be formed [166]. It is postulated [168] that this accounts for the high reactivity of carbon with oxygen. Since this gasification is conducted at comparatively low temperatures, the stable oxygen surface complex levels will be high and thus the surface will be highly reactive. The carbon dioxide gasification work was extended [166] to include additional activated carbons, Saran char, Spherocarb, a phenol-formaldehyde resin char, Ambersorb, a PVC char, two types of carbon fibre and a carbon molecular sieve.

Despite the allegations by Moulijn and Kapteijn [168] that this “fully unified” model is applicable to any oxygen-containing gas, i.e. carbon dioxide, carbon monoxide, nitrous oxide, water and oxygen, it is never experimentally verified for oxygen. However, it should be mentioned that the mechanism is strikingly similar to the one suggested by Zhuang *et al.* [160]. The notable exception exists that

Zhuang *et al.* identify the fast carbon monoxide formation step with a direct reaction between the existing semiquinone group and oxygen, without the requirement of the exotic, intermediate, off-plane oxygen complex suggested by Kapteijn *et al.*

Somewhat paradoxically, it is suggested [166] that the idea of using the ASA to determine the relative gasification activity of carbons is flawed. This assertion is based on the fact that surface complex desorption is temperature dependent, i.e. the number of complexes that are desorbable depends on the temperature (due to surface site heterogeneity) and the partial pressure of the oxidising gas. For this reason it is speculated that the total number of active sites available for reaction will depend on the temperature, due to surface heterogeneity. This suggestion is not experimentally verified but is nonetheless supported, in part, by the work of Marchon *et al.* [154]. As mentioned earlier, they showed that the TPD spectra are dependent on the temperature at which the graphite was exposed to oxygen. This indicates that the type of surface complex formed depends on its temperature of formation. Unfortunately, Marchon *et al.* did not calculate the ASAs based on these different runs.

A better formulation of Kapteijn *et al.*'s assertion may therefore be stated as: different surface complexes are formed at different temperatures due to inherent surface heterogeneity, thus it is inevitable that differences in surface heterogeneity for different carbons will yield varying estimates of the ASA, based on TPD experiments. This does unfortunately imply that unless the exact surface structures of the desorbed species under all conditions and for all carbons are explicitly known, it is not possible to infer the ASA from TPD. The influence of the formation of different complexes on the calculated ASA may well be small, but the relative surface coverage of the different complexes and the dependence of this on the gas atmosphere may not be small.

The work based on this new surface complex has also been extended to include the action of metallic catalysts [168];[169]. These are simulated using molecular orbital calculations for metallic atoms attached to the edge oxygen atom. It is postulated that metal facilitates carbon dioxide chemisorption via increased charge density on the adjacent carbon atom, which increases the amount of in-plane surface complex. However, it is concluded that this effect cannot be significant for the carbon-oxygen reaction since most of the ASA is

already covered with surface complexes. This would imply the existence of an entirely different mechanism for catalysis in the carbon-oxygen reaction.

2.9 Molecular simulation of surface complex structures

Recently, there have been an ever-increasing number of publications on the molecular orbital simulation of carbon-oxygen surface complexes [170];[171];[172];[173];[174];[175];[176];[177]. However, the final point in the previous section alludes to one of the problems with this approach: it is easy to develop a simulation that replicates one effect, but much more challenging to invent a molecular simulation capable of representing the true, fully comprehensive behaviour.

The reason is simple: there is an almost unlimited number of causes for the possible configurations of carbon and oxygen atoms at the edge of the graphitic plane. This opens a veritable Pandora's box of possible simulations to fit specific portions of experimental data, without providing any clear insight into the true scenario. Because there is no *a priori* method for selecting between the possibilities, it is inevitable that configurations will be found to fit certain aspects of the data but not others.

The danger here lies in the attitude of some researchers, e.g. [178], that these tools have the capability to solve problems that are "beyond the ability of experiments". If theoretical calculations cannot be verified experimentally, they are of little value and remain nothing more than speculation at best. Nonetheless, it is worthwhile to examine some of the possible directions suggested by this work to develop new experimental techniques that might verify these theoretical postulates.

Quantum mechanics provides a potential method for the complete description of chemical systems, including their properties and reactivities. However, such *ab initio* (from the first principles of quantum mechanics) models have encountered computational difficulties due to their severe computing time, memory and disc space requirements. In recent times, the development of new conceptual frameworks and computational methods has made it possible to overcome these difficulties [170]. The most popular of such new *ab initio* molecular orbital calculation packages is "Gaussian".

However, despite the advances in this field, it still remains necessary to select a suitable calculation level and model for the system to be studied. Computational requirements are usually determined by the choice of model chemistry, while the model that delivers a suitable description for the system under consideration depends heavily on the required outputs. The model chemistry encompasses two elements [170], namely the method and basis set. Here the method deals with the correct selection of the Hamiltonian operator (Eq. 2.2) and the basis set with the particulars of the wave function in the Schrödinger equation (Eq. 2.1). Higher model chemistry delivers more accurate results, but the computational costs increase logarithmically.

The choice of a representative molecular model can be difficult and is usually done on a trial-and-error basis, the validity of different choices being verified by some experimentally available parameters, e.g. known bond lengths and angles. An example of such an evaluation for the graphite system is given by Chen and Yang [170], who show that a polycyclic aromatic compound of no fewer than seven aromatic rings may be used to represent the molecular structure of graphite satisfactorily.

To illustrate the variability of results generated by such investigations, subsequent simulation work by Chen and Yang [179] on the exotic, off-plane oxygen complex proposed in the previous section indicates that the intermediate exists only in the carbon-oxygen reaction and is not present in the carbon-carbon dioxide reaction, in complete contradiction to the entire body of work done previously. A later study by the same researchers [173] yet again calls into question the conclusions generated by their earlier work [170], indicating that a newly considered surface configuration, so-called “o-quinone”, which is a double quinone on an armchair site, is responsible for a 30% reduction in the postulated activation energy.

The validity of the off-plane oxygen complex, referred to as the “epoxy oxygen group”, is evaluated by Montoya *et al.* [172]. They indicate that the desorption activation energy of carbon monoxide from such a structure may be as low as 8 kcal/mol, making it potentially too low to observe experimentally. Furthermore, due to the deformation of the planar graphene structure, such a group would be stable only at high temperature, making its existence very questionable. Montoya *et al.* [180] also point out some of the complexities when

using different model chemistry methods. They indicate that the popular Hartree-Fock method may be susceptible to the effect of spin contamination, leading to large errors in the calculations.

Their own work [181] indicates that the desorption energy of carbon monoxide from both carbonyl and semiquinone groups is dependent on the surrounding chemical groups. This may be a plausible explanation for the distribution of desorption energies found in TPD experiments. Furthermore, this group attempts to link *ab initio* molecular simulations to measured kinetic rate data [182] using transition-state theory. The calculated activation energy is around 400 kJ/mol, which is roughly twice the value stated previously for a wide variety of carbons [102].

Recent simulations by Espinal *et al.* [174] indicates that a reversible conversion process exists between carbonyl and ether surface groups, as suspected in the previous section [154];[160]. The interesting implication is that such an interconversion would result in the reduction of the number of active sites available for reaction and hence the reaction rate. An important potential consequence of molecular simulations is the generation of a reaction scheme that continually renews the active surface, allowing continuous, repeatable reaction. Based on a turnover model, Backreedy *et al.* [171] developed such a scheme, but unfortunately the scheme operates via dangling carbon bonds, so-called “superreactive sites”. Experimental verification of such highly unstable intermediates would be very difficult.

It is ironic that a large percentage of the work discussed here is aimed mainly at describing the gasification of amorphous, char materials, yet all the models utilise the graphite edge structure as the core of the model. A recent study [183] points out that large variations in the relative energies of the spin states occur when the cluster sizes are increased and more active sites are added. This would seem to indicate that, at a minimum, future studies should also consider the effect of amorphous structures, as well as adding graphitic layers above and below the modelled edge structure.

Another factor that it is important to consider when designing the molecular orbital model is the mechanism of adsorption. Some studies [176];[184] indicate that the energy required for dissociation of the oxygen adsorbed on graphite edge sites is extremely low. Similar investigations have been done to

explore the adsorption of carbon dioxide [185]. An extensive study on the gasification of graphite by oxygen has been done by Sendt and Haynes on both armchair [186] and zig-zag edge sites [187]. This study covers both adsorption and desorption but, interestingly, also includes possible surface migrations and rearrangements. The study was extended [188] to include rearrangement to lactone and ketene structures to account for carbon dioxide desorption.

The work on both armchair and zig-zag sites indicates that several pathways with an overall lower-than-expected energy for direct carbon monoxide formation from oxygen exist. These pathways occur via migration and surface oxide rearrangement, thus providing a qualitative insight into the existence of two distinct types of adsorption behaviour. At low-temperature stabilisation, retention of the oxygen surface complexes will be more likely, while at higher temperatures a pathway whereby one oxygen atom is retained and one is desorbed as carbon monoxide is more likely.

This is in accordance with prior investigations where secondary reactions with existing surface complexes were thought to contribute significantly to the gasification rate. However, this does significantly complicate considerations of the possible surface pathways for the carbon-oxygen reaction. It indicates that the temperature-dependent behaviour of surface complexes may be due to rearrangement and migration, rather than the discrete formation of different stable complexes. This and other work [175] also reaffirm the statement that the behaviour of a given surface complex is influenced by that of its neighbours. However, it is stated that this effect is further augmented by the fact that the rearrangement and surface migration reactions occur within a much shorter time than the desorption, thus providing further insight into the experimentally observed broad activation energy profiles for TPD.

One troubling question that remains relates to the nature of cleaned graphite edges. When a given graphite sample is heated under inert conditions to sufficiently high temperatures, it is presumed that all the oxygen-based surface complexes are removed. Usually, the samples are cooled under inert conditions and it is assumed that there is very little interaction between the inert gas and the now-exposed graphite edge sites. However, very few studies have investigated the nature of these free edge sites. This is due in part to the fact that once gasification has started, the edges are inevitably covered in surface complexes.


However, the possibilities of partially annealed edge sites or highly reactive dangling carbon bonds have been postulated for some time [148]. Such edge sites would indeed affect surface reactivity and potentially ASA development.

Recently, however, it has been proposed [60] that, under ambient conditions, these oxygen-free sites are not highly reactive free radicals but rather stable carbene-like (zig-zag) or carbyne-like (armchair) arrangements. This, on the other hand, does have implications for the oxidation reaction since these potentially stable edge sites might now be considered as reactionary intermediates. Certain more radical views of the free edge sites have been proposed [177], whereby stabilisation occurs via thermal annealing and the formation of five-membered carbon rings at the edge sites, leading to a reduction in the number of active sites and a dynamic bending and flexing of the carbon lamellae during oxidation.

The work reviewed in this and the previous section illustrates the significant complications and stumbling blocks facing the development of a truly representative molecular mechanism for the carbon-oxygen reaction. In the next section the kinetic mechanisms that have actively been applied to experimental investigations will be considered.

2.10 Kinetic mechanisms for the oxidation of graphite

Some of the earliest studies on the oxidation of graphite were conducted by Gulbransen [85];[189];[190]. He measured an activation energy for the oxidation of synthetic graphite of around 150 kJ/mol. Despite assertions that the reaction is first order at high pressure and decays to zero order at low pressure, evaluation of the reaction rate against pressure shows no such transition but a first-order dependence across all pressures tested. There is some confusion though regarding the assertion that no surface oxide is formed; this in fact refers to the formation of a protective oxide layer (such as that formed in metals) that inhibits oxidation and not to the oxygen complexes on the graphite edge sites which are observed. The first basic formulation of the reaction mechanism is also given:








Thus the steps are given as adsorption, chemical reaction and desorption. Here C_f is a bare carbon active site, $C(O)$ is a dissociatively chemisorbed oxygen and $C(CO)$ is a chemisorbed carbon monoxide, which is assumed to be the primary product. Based on the observed experimental data, it is concluded that either mobile adsorption or dissociative immobile adsorption is the rate-limiting step.

A similar study on synthetic graphite was conducted by Blyholder and Eyring [94];[191]. They measured the activation energy below 800 °C as 335 kJ/mol and the reaction order as zero. Based on this, they concluded that the desorption step is rate limiting. In addition, they measured the activation energy above 800 °C as 176 kJ/mol and the reaction order as 0.5, on the basis of which they concluded that this is the region of pore-diffusion control. Using thicker samples caused a similar transition. A subsequent study by Gulbransen and co-workers [192] reaffirmed the 800 °C transition point but showed that it depends on the sample's surface area. The reaction order of one was shown to become fractional (0.32) in the diffusion-controlled region with an activation energy of 16 kJ/mol.

An extensive study on the changes in the thermochemical properties of the reactant gases by Ong [193] indicates that the concept of reaction order when applied to gas-solid reaction kinetics is of limited utility, as is the use of Arrhenius plots to determine the activation energy, due to variations in surface coverage. The conclusions are based on the following mechanism:



One of the earliest electron microscope studies [194] revealed several unique aspects to the oxidation of cleaved, single-crystal graphite samples. When these graphites are oxidised, the activation energy is given as 268 kJ/mol but when the edges are decorated with gold to aid observation, the activation energy drops to 148 kJ/mol. Furthermore, a difference in the edge-recession rate for single versus multilayer edge pits is observed. Single-layer edge recession is 100 times slower, indicating different edge complex structures with different reactivities on mono, as opposed to multilayer edges.

Recently, multiple studies [195];[196];[197] have focused on the effects of oxidation on HOPG. In this case oxidation proceeds mainly from lattice defects and results in the formation of pits in the basal surface. These pits show a wide variety of shapes and configurations [196], depending in part on the type of defect where the pit formed. The pitting behaviour has also been shown to be temperature dependent [195] and, as mentioned earlier, this was shown by Thomas [10] to be related to reactivity differences between the graphite armchair and zig-zag edges.




Since it is common practice to use an energetic gas plasma to clean microscopy samples, it is interesting to note that Paredes *et al.* [197] found that this induced significant surface roughening. When HOPG samples exposed to air and oxygen plasma respectively were compared, the surfaces exposed to air were found to be flat on the atomic level and oxidation proceeded only at existing defects in the lattice. In contrast, the samples exposed to the oxygen plasma showed a uniformly roughened topography induced by continuous creation of defects in the basal plane. This indicates that caution should be exercised when plasma cleaned samples are examined.



Although most of these studies provide good topographic information on the oxidised surface, only a few have focused on direct measurement of the kinetic parameters for oxidation. Pit growth analysis can be used to obtain several kinetic parameters during oxidation [198]. Furthermore, a widely used technique for controlling and manipulating the amount of defects in the HOPG structure is via ion bombardment [199]. Using this technique enables the number of both mono- and multilayer pits to be controlled and their relative growth rates determined.

Several studies have utilised this technique [199];[200];[201];[202] to prove the assertion by Evans and co-workers that monolayer edge recession is slower than multilayer, but with lower ratios than those reported by Evans. Stevens *et al.* [200] determined the activation energy for edge recession in mono- and a variety of multilayer edge pits to be between 165 and 180 kJ/mol. They also attempted to measure the reaction order in oxygen but the results varied widely. They postulate that this can be explained by oxygen surface complex interactions between adjacent graphene layers, a fact substantiated by the observation that the etch rate increases markedly only for very shallow pits.

Hahn *et al.* also studied the growth rates of etch pits on HOPG [201];[202] and concluded that the activation energy ranges from 96 to 130 kJ/mol at low temperatures, depending on the depth of the pit, and is 143 kJ/mol for etching at higher temperatures. Due to difficulty in estimating the exact depth of the pit, absolute rate calculations are uncertain but nonetheless the studies by both Hahn *et al.* and Stevens *et al.* deliver comparable results for the pre-exponential factor. The units are given as the number of carbon atoms removed per unit time per number of edge carbon atoms and the values are: $20\text{--}300 \times 10^{10} \text{ (s}^{-1}\text{)}$ [200] and $1 \times 10^{10} \text{ (s}^{-1}\text{)}$ [202]. Assuming each carbon atom occupies an area of 0.083 nm^2 , these rates can be converted to: $48\text{--}720 \times 10^6$ [200] and 2.4×10^6 [202], with units of $\text{g carbon reacted.m}^{-2} \text{ ASA.s}^{-1}$.

The activation energy found by Barton *et al.* [147] for the desorption of carbon monoxide from synthetic graphite is very similar (276 kJ/mol) to that found by Evans *et al.* [194], indicating that for these samples desorption was rate limiting. The study by Ahmed and Back [203] on pyrolytic graphite films also measured a similar activation energy of 290 kJ/mol for the desorption of carbon monoxide by TPD and was one of the earlier models to include the concept of different surface complexes. The kinetic mechanism used was:





where $C(O_2)$ is an adsorbed oxygen before conversion to the stable surface oxide (CO) .

Importantly, both carbon monoxide and carbon dioxide formation are secondary with respect to the formation of the stable complex. Reactions (2.24)–(2.26) represent mechanisms other than simple desorption whereby the surface complex contributes to carbon dioxide formation. This mechanism implies that the surface reactions become less important relative to the desorption at low oxygen pressure, but even at high pressure the rate will still depend on the oxygen concentration. Since the stable complex acts as an intermediate, a small amount of this complex can lead to a large amount of product formation and, in a sense, the complex acts as a catalyst.

Ahmed and Back suggest that this indicates that the ASA may not be a direct indicator of the reactivity and instead suggest that the turnover number, defined as reaction rate (in $g/m^2/s$) divided by ASA (m^2/g), normalises the reaction rate. It is important to note that the turnover number still has units of grams of carbon reacted per grams of carbon remaining at a given time per second. Thus this ASA value is in terms of the grams of carbon remaining at a given time. This is only possible due to the given reaction mechanism through which the product gases are always produced via the same surface complex; each reaction to produce them leads to a regeneration of the complex. The turnover number can be converted into the rate measured on HOPG shown earlier through division by the Arrhenius exponential, which gives a value of $0.1 \times 10^{10} (s^{-1})$, in reasonable agreement with Hahn *et al.* [202].

Unfortunately, when this approach was applied to a wide variety of published data, the results were inconclusive. Despite the agreement of the carbon monoxide desorption activation energy with the previous studies, the activation energy for the overall gasification rate was far less than this (170 kJ/mol). It is concluded that the gasification proceeds mainly via reactions

(2.24)–(2.26) and that, due to the low reaction temperature and pressure, the stable complex surface coverage is reduced.

A similar activation energy (186 kJ/mol) was measured by Özgen and Rand [204] for a type of natural flake graphite with a relatively high impurity content of 10% across a temperature range of 700–900 °C. A much lower activation energy (117 kJ/mol) is reported by Cascarini de Torre *et al.* [205] for a graphitised carbon black and the transition to diffusion control was shown to occur at much lower temperatures (400 °C). However, their study indicates that the surface area of the material increases continuously across the conversion range of 5–75% during the kinetically controlled oxidation, which seems to imply that the material has very large amounts of closed-pore volume. This makes the determination of activation energy dubious as the reaction rate is a combination of kinetic activation and deterioration of the pore wall structure, i.e. geometry and thickness, as well as crystallite orientation in the wall.

Some studies have found that normalising to the ASA is adequate for representing a variety of experimentally obtained oxidation data. Zaghbi *et al.* [206] determined the ASA from crystallite calculations based on XRD and Raman spectra. They obtained an activation energy of 188 kJ/mol for a collection of natural graphite samples and 170 kJ/mol for a synthetic graphite. This is based on the simple rate expression developed by Radovic *et al.* [149], mentioned in Section 2.7:

$$R_R = \frac{-1}{m_t} \frac{dm_t}{dt} = \frac{1}{1-\alpha} \frac{d\alpha}{dt} = k_f C_f \quad (2.27)$$

where R_R is the dimensionless reaction rate, m_t is the sample mass at time t and α is the dimensionless degree of conversion, which may be defined as: $(m_0 - m_t)/m_0$; where m_0 is the sample mass at the start of the experiment. Also k_f is the reaction rate constant and C_f is the concentration of carbon active sites, which is assumed by Zaghbi *et al.* to be the fraction of edge sites times the BET surface area.

From their data the turnover number can be calculated as $2.2 \times 10^{10} \text{ (s}^{-1}\text{)}$, giving excellent agreement with the HOPG value. One point to bear in mind is that this result is based on the determination of the reaction rate constant from

the initial slope of the carbon weight loss curve during an isothermal experiment (i.e. $\alpha \rightarrow 0$) in Eq. (2.27). Thus the study gives no data regarding the development of ASA as a function of conversion and allows reactivity comparisons based on initial conditions only.

A similar study was carried out by Kim *et al.* [207] on nuclear-grade synthetic graphite IG-110. In this case the reaction temperature was chosen to be very low and therefore mass loss was very slow. This made it straightforward to determine the initial reaction rate. Interestingly, by comparing samples with a known but varied external surface area they were able to determine the internal surface area. The activation energy was given from a previous publication to be 215 kJ/mol and the turnover number can be estimated as $1.6 \times 10^{10} \text{ (s}^{-1}\text{)}$.

Fuller and Okoh [100] described the kinetic behaviour of IG-110 across the entire conversion range. They used the random pore model mentioned earlier to model the surface area development, thus assuming a constant ratio of ASA to TSA. Despite concluding that the experiments were conducted in the Zone II region, they asserted that the samples were under kinetic control and measured an activation energy of 188 kJ/mol. The experiments were conducted between 450 and 750 °C and the turnover number may be estimated as $0.2 \times 10^{10} \text{ (s}^{-1}\text{)}$.

The previous studies highlight an important trend in oxidation studies on graphite. In general, these studies have been aimed at understanding the properties of synthetic graphite materials. This is due to the fact that these are the only bulk forms of graphite and are therefore of engineering importance in a variety of applications. Furthermore, these studies tend to focus on determining the bulk properties. However, as pointed out in previous sections, these are highly porous materials and subject to mass transfer limitations. The previous group of researchers circumvented this challenge by incorporating a model for ASA development and the groups prior to that by using only initial rates to evaluate the properties of different samples.

In certain cases, however, it becomes important also to understand the high-temperature behaviour. Xiaowei *et al.* [208] showed that for IG-11, a material closely related to the IG-110 used in the previous study, the transition to diffusion control occurs around 600 °C. Below this temperature the activation energy was measured at around 159 kJ/mol and above this it dropped to 72

kJ/mol, indicating a transition from Zone I to Zone II. Above 800 °C the sample showed very little activation energy and the reaction was controlled by the boundary layer, consistent with Zone III behaviour. This material contained 479 ppm ash and the initial turnover number at 700 °C may be roughly estimated as $0.5 \times 10^{10} \text{ (s}^{-1}\text{)}$.

However, several studies [209];[210];[211] tend to focus exclusively on the diffusion-limited region. Unfortunately, these studies contain very little experimental kinetic data for comparison with other studies. In some cases [112];[113];[212];[213] the simulations to describe these conditions can become very complex, especially when it becomes necessary to model the flow around stationary graphite samples using computational flow dynamics.

To describe these conditions accurately, the intrinsic kinetics are needed, particularly in Zone II where a combination of diffusion and reaction is taking place [213]. Furthermore, these situations depend heavily on either the measurement of the effective diffusion coefficient [210] or structural knowledge of the complex pore evolution of graphite. All of these factors make modelling the oxidative behaviour in this region very difficult without clear knowledge of the transition regions and pore evolution as a function of conversion.

One method mentioned earlier for circumventing this limitation is to conduct the kinetic experiments in a fluidised bed reactor. However, this approach, too, has its disadvantages. Bews *et al.* [89] conducted an extensive study on the oxidation of synthetic graphite particles in a bed fluidised with sand. They found the activation energy of the oxidation to be 179 kJ/mol and the reaction order around 0.5. The reaction rate itself was measured in a peculiar way. The rate is based on the total mass of carbon fed to the bed, the initial BET surface area and the oxygen concentration. This in no way takes into account the variation of active surface area with conversion, and calculates an average rate across the full conversion range. The calculated turnover number is therefore quite low: $0.036 \times 10^{10} \text{ (s}^{-1}\text{)}$.

A fairly comprehensive review of the past 30 years' worth of literature (at the time) on the carbon-oxygen kinetic mechanisms was compiled by Hurt and Calo [214]. They found significant evidence to suggest a reaction order for the Zone I regime between 0.6 and 1.0, and the activation energy between 105 and 180 kJ/mol with a high proportion between 130 and 150 kJ/mol. It should be

noted that this is for any form of carbon, from graphite to amorphous char and soot. Four kinetic models were evaluated in terms of their predictive accuracy and ability to model the observed transitions in an effort to find a semi-global intrinsic kinetic model.

The first model evaluated is the *empirical global power law model*. This is very similar to Eq. (2.27), with the notable exception that it is not linked to nor dependent on the ASA.



$$R_R = kP_{O_2}^n \quad (2.28)$$

Due to its inability to model the transition from Zone I to Zone II, it is disregarded as a potential semi-global model.

The second possibility is the *non-dissociative form of the Langmuir-Hinshelwood expression* where there is no surface reaction between adsorbed species:

$$R_R = \frac{k_1 k_2 P_{O_2}}{k_1 P_{O_2} + k_2} \quad (2.29)$$


The reaction steps are given as:



This mechanism is the classic description of the competition between adsorption and desorption. In general, the activation energies for desorption are much higher than those for adsorption [215]. Thus the Langmuir-Hinshelwood expression exhibits two limiting cases: first-order adsorption control at high temperature and zero-th order desorption control in the low-temperature limit. However, as mentioned in previous sections, this behaviour is significantly complicated by the behaviour of the surface coverage of the reactionary intermediates. In fact, exactly the opposite behaviour is observed for carbon materials [216]. This negates any mechanistic significance of this expression.

Furthermore, since the model is unable to account for the observed persistent power law behaviour [217], it is discounted as a general kinetic model.

The studies mentioned earlier [160];[165] that utilised isotopically labelled atoms together with, at the time, more recent studies using this approach [218], provide direct evidence for the reactionary pathway via a previously formed surface complex. This point led to the suggestion of the third possibility, i.e. the following *three-step semi-global mechanism*:





which is a combination of Eqs (2.21–2.24) as suggested by Ahmed and Back [203] and is a reduction of the qualitative mechanism suggested by Zhuang *et al.* [160], with the corresponding reaction rate law:

$$R_R = \frac{k_1 k_2 P_{O_2}^2 + k_1 k_3 P_{O_2}}{k_1 P_{O_2} + k_3 / 2} \quad (2.35)$$


It is interesting to note that initially the reaction rate expressions are given in terms of the fraction of active sites covered by the surface complex. However, when the final rate expressions are given, this dependence is neglected. Using a fairly empirical choice of activation energies based on literature values, it is shown that this expression is capable of representing the wide variety of reaction order data that has been collected. Specifically, the model predicts the transition from high to low order as temperature is increased, but allows absorption control at very high temperatures (which a simple inverted Langmuir-Hinshelwood law would not). Furthermore, the model is also shown to provide an estimate of the “effective activation energy” for the CO:CO₂ ratio found in the literature.

The critical element in this model is the inclusion of Eq. (2.33), i.e. the carbon dioxide-producing step via the “stable” surface complex. However, this model is still not capable of fully accounting for a few known experimentally observed features of the reaction, including the production of carbon dioxide as a

second product of thermal desorption, site heterogeneity, increasing oxide surface density with an increase in temperature, constant n th-order behaviour over two orders of magnitude of oxygen partial pressure and different reaction orders for carbon monoxide and carbon dioxide. The following expansion to the previous model is therefore proposed:



In this case an additional type of active site, C_b , with a higher activation energy is added. This model is shown to be consistent with the observed persistent n th-order behaviour. In its general form, the rate expression is very complex and this model was not treated in any amount of detail. In 2001 the concept of a turnover model was revived by Haynes [219]. The model is very similar to Eqs (2.36–2.41), but a stochastic approach is used to characterise the heterogeneity of the surface sites in order to develop a turnover-based model. A slight modification to reaction (2.37) is proposed based on TPD and the literature data:



Unfortunately, it was chosen to focus only on carbon monoxide formation and ignore carbon dioxide. The model uses a coverage-dependent adsorption energy and the desorption rate is determined by integration across a distribution of sites with varying desorption activation energies. Based on TPD experiments with variable soak time, a Gaussian distribution is used. The model shows a fair correlation for the gasification rates of Spherocarb at three different

temperatures. The experimentally observed activation energy of $115 \text{ kJ}\cdot\text{mol}^{-1}$ is correctly predicted at an oxygen pressure of 2.7 Pa. The work was further extended by Hurt and Haynes [95] who used a pseudo-steady state version of the model to show that it is capable of describing the persistent power law behaviour mentioned previously, as well as the reaction orders, activation energies and reaction rates for a wide variety of coal and polymer chars. It also predicts the low reaction order observed for carbon black.

Recently, what could be considered as a culmination of this approach was applied to a synthetic furfuryl alcohol polymer char by Campbell and Mitchell [220]. The model incorporates no fewer than 15 possible reaction pathways, including physical adsorption, chemisorption, surface migration and desorption. However, it does not include a direct oxygen surface complex reaction step. It also incorporates the random pore model for predicting surface area development, and a symmetrical fourth-order polynomial desorption energy distribution is used for both carbon monoxide and carbon dioxide. The model's prediction of the observed oxidation behaviour over $200 \text{ }^\circ\text{C}$ is reasonable, but perhaps not unexpected given the number of adjustable parameters it contains. The authors conclude that a minimum of four types of active site are needed to account for the behaviour and that a failure to link the surface area evolution during oxidation with surface species concentration can lead to significant errors.

These approaches seem capable of explaining the observed behaviours and providing insight into the complexity of the wide and varied observed behaviours. However, the exceedingly large task of characterising all possible surface complexes for every carbon material of technological significance makes them highly impractical [99]. Furthermore, the number of parameters in these models may become excessive, allowing apparent goodness of fit but providing little applicability outside of the tested range. Based on these complex formulations, a variety of models and experimental approaches have been proposed.

Using a fairly generic reaction mechanism and the concept of stable and unstable oxide formation having a distribution of desorption energies, Battye and Ashman [221] developed a four-step reaction mechanism which they applied to a phenol-formaldehyde polymer char. They came to a new set of conclusions including:

- Surface complexes are dependent on conversion.
- The formation of carbon dioxide proceeds mainly via surface decomposition.
- The reaction rate is initially high due to the availability of vacant sites, showing a decrease as these sites are filled, followed by an increase as carbon dioxide-yielding complexes are formed.

Another approach is suggested by Li and Brown [222], who utilise temperature-programmed oxidation. During such an experiment the sample is subjected to a temperature ramp which increases at a fixed rate. Based on work using graphite, charcoal and coke, they conclude that oxide formation retards oxidation for all carbons.

Many of the conclusions in these studies are in direct contradiction to previous work. Furthermore, very few consider the development of active surface area concurrently with the measured surface complex formation. In some cases the experimental methodology is unclear and the validity of the assumptions made is difficult to assess. These facts serve to increase the uncertainty surrounding conclusive proof of the accuracy and reliability of the complex kinetic mechanisms that have been discussed. Thus a clear and conclusive reaction pathway is still contentious.

A factor that may also contribute to the seemingly contradictory results obtained in the recent literature is the complexity of the measured TGA data. The oxidative reactivity of even slightly different carbon materials can be dramatic [223]. Moreover, synthetic graphite is a combination of various precursors with varying degrees of graphitisation. This fact is highlighted by the study by Moorman *et al.* [224] who utilised the composite nature of the TGA response curve to distinguish clearly between the binder and the filler material. This approach even allowed the determination of different activation energies for these materials and showed the pronounced temperature dependence of the shape of the composite reaction rate curve. This has also been illustrated for other carbon materials to discern between mixtures of coke and coal char [225]. Thus the comparison of complex reaction mechanisms based on materials with differing heterogeneous surface characteristics should be undertaken with care.

Recently, an alternative approach seems to have become prevalent [226];[227];[228] whereby experimental data are simply fitted using completely empirical models. These models are generally based on conversion versus time plots, utilising different polynomial expressions for time, reduced time (time divided by the time required to reach 50% conversion) or conversion. A unique approach [229] compares the number of “pinholes” formed during oxidation for 10 min at 950 °C as a measure of reactivity, irrespective of depth or structure. Although these “models” may provide adequate descriptions for the behaviour of specific graphite samples, extrapolation of the results to regions outside of the test space would be perilous. Also, meaningful comparison of samples from different origins is impossible since there is no link to measurable characteristics such as ASA or impurity content.

The semi-consistent experimental results, especially for graphite, suggest that normalisation within a coherent framework must be possible to some extent. The extensive variations in the mechanistic studies, however, perhaps indicate that it is time for a new approach to modelling the oxidation of graphite. It is pertinent to note that very few studies have focused on first understanding the underlying development of the active surface area geometry as a function of conversion. In the next section a few alternative approaches to modelling the oxidation of graphite are considered. In addition, a technique that has seen only limited utilisation in investigating the gas-carbon reaction but is widely applied to other gas-solid reactions, namely solid state kinetics, is investigated.

2.11 Other modelling techniques

Fairly recently, attempts have been made to use computer-based techniques such as Monte Carlo simulations and finite element methods to model the properties of graphite, including the oxidative behaviour. Broadly speaking, *Monte Carlo simulations* are a class of computational simulation that rely on repeated random sampling to generate their results. These techniques are used to model a wide variety of physical or chemical phenomena ranging from reaction to diffusion [230]. Kyotani *et al.* [231] combined the approach with molecular orbital theory, described earlier, to simulate carbon gasification. A variety of hexagonally based molecular models are considered and the free valence at

each carbon atom in the lattice is taken as a measure of the reaction probability. This model utilises different reactivities for armchair and zig-zag sites. It leads to a few important conclusions, namely that the specific reaction rate increases monotonically with conversion and is independent of crystallite shape. It is mentioned that the former is in agreement with commonly observed carbon gasification and the latter is due to a tendency of the ratio of armchair to zig-zag sites for any given configuration to converge to the same value as conversion proceeds. Furthermore, it is confirmed that the specific reaction rate depends on crystallite size since the ratio of edge sites to bulk sites changes.

A similar molecular model is used by Stevens and Beebe [232] to model the formation of single and multilayer pits in HOPG. However, in this case the relative reaction probability was not calculated using molecular orbital theory. Instead, a collection of different edge-site possibilities was postulated, ranging from dangling carbon atoms to sites with two or more neighbours, and different configurations for armchair and zig-zag sites were considered. A different reaction probability between zero and one was intuitively assigned to the different types of edges site. These reaction probabilities were varied and the effect on pit development was monitored. Based on this work, it was found that the formation of hexagonal, as opposed to roughly circular, etch pits is dependent only on the difference in reactivity of armchair and zig-zag sites, i.e. other sites play virtually no role in the pit shape development. This is consistent with the work of Thomas [10], as mentioned earlier, who found that the transition from hexagonal, armchair pits to hexagonal zig-zag pits occurred via intermediate semi-circular pits as the temperature was increased.

A slightly different approach was taken by Bhatia [57] to simulate the evolution of layered graphite crystallites. The hexagonal structure of graphite was discretised in a manner entirely different from that used by the two previous authors. A crystallite was assumed to be adequately represented by several concentric rings of carbon atoms. Since reaction at an edge site opens up valences in the adjacent atoms, thus increasing their reactivity, the reaction rate along the perimeter was thought to be much higher than that in the radial direction. Therefore the reaction was assumed to occur layer by layer, with all the atoms on a given circle reacting at once. Based on this structure and a few assumptions regarding the reactivity of adjacent planes, the model is extended to

allow prediction of the evolution of particle size and pore structure. The simulation is shown to be capable of predicting the evolution of the reaction rate with conversion of SpheroCarb with a high degree of accuracy.

In general, the *finite element method* is a numerical technique for finding the approximate solutions of partial differential equations. A given structure is chosen across which the equations have to be solved and the structure is subdivided into a finite number of subelements. Then the equations across these elements are approximated and solved. The choice of mesh size, i.e. the size of the individual subelements, is a critical design parameter to ensure representative results. This approach was applied by Hall *et al.* [233] to describe the fracture mechanics of nuclear graphite. In this case the microstructure of graphite is assumed to comprise three possible elements: a binder phase, a filler phase and a porous phase. Each phase was assigned its own properties and model behaviour. Using this approach it was possible to model adequately the radiation-induced mechanical property changes of the graphite. If the mesh size is reduced to the level of individual atoms, this technique is similar to those mentioned previously. However, as a bulk approximation it has not yet been applied to simulating graphite oxidation.

2.12 Solid state kinetics

Following the creation of the original thermo-balance, a variety of highly accurate thermal analysis instruments have been developed. These include thermogravimetry (TGA), discussed in Section 2.3.5.6, as well as differential thermal analysis (DTA) and differential scanning calorimetry (DSC). DTA and DSC are related in the sense that both are aimed at determining the thermal properties of a solid, e.g. the heat capacities or the latent heats associated with phase transitions. However, DTA requires assumptions regarding the heat capacity of a sample, whereas DSC does not, making it the preferred technique for these measurements. Based on these analytical techniques, a variety of computational techniques have arisen to obtain the sought-after properties from the data delivered by the equipment.

Specifically, the analytical techniques associated with obtaining the kinetic parameters of any reaction involving a solid reactant can be grouped under the

general heading of *solid state kinetics*. This approach has been applied to a multitude of gas-solid reactions, crystalolysis, solid decompositions and a variety of complex phase transitions. Hence the reaction mechanisms covered by this approach are extremely broad. Furthermore, the presence of a solid reactant brings with it significant complications, compared with, for example, gaseous reactions, since the concept of concentration is no longer valid, especially due to heterogeneity. This approach provides a comprehensive look at the mechanisms that influence the kinetics of solid reactants in general. These insights are applicable to the gas-carbon reaction but from a fresh angle. The approach is generally based on concepts from the theory of homogeneous chemical kinetics [234], and it has come under increasing doubt in recent times [235]. The starting point for the kinetic modelling is the following expression [236]:

$$R_{T0} = \frac{d\alpha}{dt} = k(T)f(\alpha) \quad (2.43)$$

where α is still the dimensionless degree of conversion but in this case defined more accurately as $(m_0 - m_t) / (m_0 - m_\infty)$; where m_∞ is the sample mass at the end of the experiment.

In general, the graphite samples have very little ash but when ash is present, this expression compensates for it. Then $k(T)$ is the temperature-dependent rate constant and $f(\alpha)$ is the so-called “conversion function” or reaction model and represents the underlying mechanism governing the reaction. This expression is very similar to Eq. (2.27) which is widely applied to gas-carbon reactions. The issue at hand is how valid is Eq. (2.43) since it is based simply on the analogous expression for homogenous kinetics [104]:

$$-\frac{1}{V} \frac{dN_A}{dt} = k(T)f(C_{A\dots n}) \quad (2.44)$$

where V is the reactor volume, N_A is the number of moles of the reactant under consideration, $k(T)$ is the temperature-dependent rate constant and $f(C_{A\dots n})$ is some function of the concentration of all species participating in the reaction. This expression was first developed [234] using collision theory and later as transition

state theory (TST). The TST model is based on the assumption that some reaction rate-controlling intermediate follows the high-energy collision of freely moving reactant species. This unstable intermediate undergoes some bond redistribution and then decays into stable molecular fragments. Most importantly, as first suggested by the great Swedish chemist Arrhenius, the energy distribution of gases obeys the Maxwell-Boltzmann distribution. This allowed the definition of the temperature-dependent reaction rate constant as follows:

$$k(T) = k_0 e^{-\frac{E_A}{RT}} \quad (2.45)$$

where k_0 , sometimes given as A , is the pre-exponential factor or frequency factor, E_A is the activation energy of the reaction, R is the gas constant ($8.314 \text{ J}\cdot\text{mol}^{-1}\cdot\text{K}^{-1}$) and T is the temperature in Kelvin.

The activation energy is considered to be the energy barrier to the bond-redistribution step mentioned earlier. Several key differences exist when considering the reaction of a gas with a solid however. The solid precursor species retains links with its neighbours after the initial collision and the energy distribution in solids differs from the Maxwell-Boltzmann behaviour. Thus technically, the model cannot be directly applied to these reactions because the nature of the intermediate is difficult to isolate and the role of the energy distribution in activating the chemical change is less certain [234].

Furthermore, the reaction mechanism almost always involves multiple steps, as mentioned previously. For the gas-carbon reaction, this possibly includes adsorption, dissociation, chemisorption, migration, transition, surface reaction and desorption. For crystal growth reactions, even more steps are of concern, e.g. melting, decomposition, nucleation, growth and interfacial phenomena. Transitions between different rate-controlling steps can occur and all of the steps can certainly not be thought of as obeying homogeneous chemical kinetics. This directly questions the validity of the assumption that the kinetic observations alone contain the necessary information from which the rate-controlling steps can be deduced [237]. Despite this fact, these expressions are applied to an ever-increasing collection of kinetic data for an ever-widening sphere of solid reactions. These studies focus solely on the kinetic characteristics

of the reaction and very little effort is made to link these to clear reaction mechanisms [238]. The probable cause for this is the ease with which large amounts of thermal analysis data can be accumulated. Hence the focus was on developing an automated scheme to collate these data into largely empirical formulations which are capable of representing all the observations. However, this has resulted in a lack of coherence within this large accumulation of data without a consistent framework from which a generally applicable, structured, mechanistic understanding can grow [238]. This conclusion seems to bear a resounding similarity to the ever-increasing collection of gas-carbon kinetic data.

Currently, a general formulation of the solid state kinetic approach may be summarised as follows [239]. First, an experimental methodology is chosen. An experimental formulation can include purely isothermal or non-isothermal steps. A non-isothermal methodology usually involves imposing a temperature ramp on the reactants at a constant rate, e.g. β °C.min⁻¹. This approach has the advantage that the gas atmosphere, as determined by the purge gas connected to the TGA, is constant throughout the experiment. This is in contrast to isothermal experiments which, depending on the nature of the reaction, are usually kept under inert conditions until the reaction temperature is reached, at which point the purge gas flow is changed to the relevant reactive gas. The dynamic nature of this gas change in the reactor must be taken into account. A more general formulation of Eq. (2.43) may be given as:

$$\frac{d\alpha}{dt} = k(T)f(\alpha)a(\alpha, T, P, \dots) \quad (2.46)$$

where $a(\alpha, T, P, \dots)$ is the accommodation function. This function is understood as accommodating the actual reaction conditions in the fundamental kinetic equation [240].

An example of such a function would be the power law behaviour of the pressure in Eq. (2.28), which takes the effect of pressure on the reaction into account. The experimental data are collected, usually in the form of a mass versus time or temperature signal. The mass signal can be easily converted to conversion by dividing by the starting mass and can then be differentiated with respect to time to give the left-hand side of Eq. (2.46), i.e. the dimensionless

reaction rate. In non-isothermal experiments, it is sometimes customary to differentiate with respect to temperature, but since the temperature is just some constant (i.e. the ramp rate, β) multiplied by the time, this is analogous.

Assuming the reaction is truly under the control of the kinetic mechanism underlying the reaction model, the accommodation function may be set to unity. If an isothermal experiment was conducted, the reaction rate constant will have some fixed value and a plot of the dimensionless reaction rate against the conversion will directly represent the reaction model. A wide variety of mechanisms have been developed based on semi-empirical models, e.g. the equations of Avrami-Erofeev [241];[242];[243];[244], Prout-Tompkins [245], Jander [246], and Ginstling-Brounshtein [247]. A summary of the most widely used of these models and a few empirical models is presented in Table 2-12. The assumptions necessary for and the derivation of these models are neatly summarised by Khawam and Flanagan [248].

Table 2-12: Kinetic mechanisms in solid state reactions [248]

Kinetics	Comments	$f(\alpha)$
Power law	Acceleratory α -time	$n\alpha^{(n-1)/n}$
Exponential	Acceleratory α -time	α
Avrami-Erofeev	Sigmoid α -time ($n = 2,3,4$)	$n(1-\alpha)[- \ln(1-\alpha)]^{(n-1)/n}$
Prout-Tompkins	Sigmoid α -time	$\alpha(1-\alpha)$
Geometric boundary control	Deceleratory α -time (disc $n = 2$, sphere $n = 3$)	$n(1-\alpha)^{(n-1)/n}$
Diffusion control one-dimensional	Deceleratory α -time	$\alpha / 2$
Diffusion control two-dimensional	Deceleratory α -time	$[- \ln(1-\alpha)]^{-1}$
Diffusion control three-dimensional	Deceleratory α -time (Jander's equation)	$\frac{3}{2}(1-\alpha)^{2/3} [1 - (1-\alpha)^{1/3}]^{-1}$
Ginstling-Brounshtein	Deceleratory α -time	$\frac{3}{2} [(1-\alpha)^{1/3} - 1]^{-1}$
"Order of reaction"	Deceleratory α -time ($n = 0,1,2,3$)	$(1-\alpha)^n$

The form of the experimentally identified conversion function is then compared with the expressions in Table 2-12 to find one with the same behaviour. Based on this, the underlying mechanism can potentially be identified. A technique for evaluating the various regions of the α -time or reaction rate-time curves is suggested by Galwey and Brown [249]. They point out that it is usual to find that the onset of the dominant rate-controlling process is usually preceded by an “induction period”. In some cases this may be caused by secondary processes or the decomposition of surface material or impurities. This region should be subtracted from the curve before kinetic analysis is performed.

By repeating the analysis at different temperatures, the pre-exponential factor and the activation energy can be found. This can be done in a variety of ways, e.g. by constructing a conventional Arrhenius plot at a given conversion. Evaluating any one of the models listed in Table 2-12 is a complex task, as pointed out by the extensive discussion on the Prout-Tompkins expression by Brown [250]. This evaluation points out the necessity for an intimate understanding of the mechanism under consideration and how subtle variations in the experimental data should lead to clear implications for the interpretation of the said mechanism. In an effort to simplify the procedure, however, a general rate expression has been proposed to represent all the models given in Table 2-12. This is known as the Sesták-Berggren equation and is entirely empirical [240]:

$$f(\alpha) = \alpha^m (1 - \alpha)^n [-\ln(1 - \alpha)]^p \quad (2.47)$$

This curve was intended to give a preliminary estimate of the reaction mechanism by fitting the curve with its three adjustable parameters (m,n,p) to the experimental data. However, due to difficulties with finding the optimal numerical solution to Eq. (2.47) and the fact that overlapping values for the exponential factors are possible, this expression cannot give a direct determination of the true mechanism [236]. This is a step backwards according to the prior discussion, as illustrated by the recent publication by Perez-Maqueda *et al.* [251]. They maintain that the kinetic mechanism can be obtained via a simplified form of Eq. (2.47) and a combination of isothermal and non-isothermal data. The problem with this approach is two-fold: firstly, the solutions are not always unique and secondly,

this assumes that only one rate-controlling process is active across the entire range of conversion studied.

Evaluating non-isothermal data is slightly more complex. It has been shown that it is not possible to determine the kinetic parameters and kinetic model from a single non-isothermal experiment [252];[253]. A variety of approaches has been suggested to overcome this, e.g. a direct modification of the kinetic rate law based on non-isothermal considerations [254]. This was based mainly on the introduction of a new concept known as “reduced time”. Firstly, one may define generalised time (τ) as:

$$\tau = \int_0^t \exp\left(-\frac{E_A}{RT}\right) dt \quad (2.48)$$

By differentiating this expression and substituting into Eq. (2.43), one obtains:

$$\frac{d\alpha}{d\tau} = k_0 f(\alpha) \quad (2.49)$$

This corresponds to the generalised kinetic equation at infinite temperature [255]. By using a predetermined, constant activation energy, the reaction rate at infinite temperature, i.e. $d\alpha/d\tau$, may be plotted against α using the experimental data and the following expression:

$$\frac{d\alpha}{d\tau} = \frac{d\alpha}{dt} \exp\left(\frac{E_A}{RT}\right) \quad (2.50)$$

At this point a so-called “master plot” can be constructed as follows. A reference conversion is chosen, e.g. $\alpha = 0.5$. Then the following expression can be derived from Eq. (2.49):

$$\frac{d\alpha/d\tau}{\left(d\alpha/d\tau\right)_{\alpha=0.5}} = \frac{f(\alpha)}{f(0.5)} \quad (2.51)$$

where $f(0.5)$ is a constant for a given kinetic model function. Thus a plot of the reduced reaction rate, i.e. the left-hand side of Eq. (2.51), is equivalent to the plot of the theoretical, reduced conversion function, i.e. the right-hand side of Eq. (2.51). Given a previously determined activation energy, this method is equivalent for isothermal and non-isothermal data since the measured temperature is utilised directly in Eq. (2.50). Hence any experimental data can be compared with a master plot to identify the mechanism. This is known as the differential form of the master plot. As an alternative, the integral form of the master plot may also be used. Rearranging Eq. (2.49) yields:

$$\frac{d\alpha}{f(\alpha)} = k_0 d\tau = k(T) dt \quad (2.52)$$

and integrating:

$$g(\alpha) = \int_0^{\alpha} \frac{d\alpha}{f(\alpha)} = k_0 \int_0^t d\tau = k_0 \tau \quad (2.53)$$

where $g(\alpha)$ is the integral form of the reaction rate, which may be used to generate a master plot using the expression:

$$\frac{\tau}{\tau_{\alpha=0.5}} = \frac{g(\alpha)}{g(0.5)} \quad (2.54)$$

In this case the reduced generalised time is plotted against the theoretical master plots of reduced integral reaction rate. For isothermal conditions Eq. (2.53) may simply be written as:

$$g(\alpha) = \int_0^{\alpha} \frac{d\alpha}{f(\alpha)} = k(T)t \quad (2.55)$$

However, for non-isothermal conditions the integral in Eq. (2.48) must be solved. For a linear heating rate of β the expression has no analytical solution. A variety of methods exist for solving the temperature integral [256], each with varying degrees of accuracy. Using a method that solves this integral with a high degree of accuracy is of paramount importance to obtaining useful kinetic data. One approach involves approximation by a convergent series. By rearranging Eq. (2.48) we may obtain:

$$\tau = \frac{1}{\beta} \int_0^T \exp\left(-\frac{E_A}{RT}\right) dT = \frac{E_A}{\beta R} \int_x^\infty \frac{e^{-x}}{x^2} dx = \frac{E_A}{\beta R} p(x) \quad (2.56)$$

where $x = E_A/RT$ and the function $p(x)$ may be approximated with an accuracy of better than $10^{-5}\%$ [257], as:

$$p(x) = \frac{e^{-x}}{x} \cdot \frac{x^3 + 18x^2 + 86x + 96}{x^4 + 20x^3 + 120x^2 + 240x + 120} \quad (2.57)$$

Thus the integral approach can be applied for both isothermal and non-isothermal conditions, although it is generally accepted that differential methods are more suitable for obtaining meaningful kinetic parameters than integral methods [258]. Furthermore, it is recommended that once a kinetic model has been selected, the invariability of activation energy with conversion is confirmed [259]. This procedure involves back-calculating the activation energy for both isothermal and non-isothermal experiments for both differential and integral methods.

There are also other non-isothermal methods, e.g. constant-rate thermal analysis [257] and sinusoidal temperature modulation [251]. These methods are non-linear and in some cases the time-temperature dependence is not known, e.g. during secondary effects. In this case it is necessary to use an empirical fitting function such as Eq. 2.47. Finally, it is also possible to combine the differential and integral techniques into a single approach. All of these methods allow the evaluation of non-isothermal data given a known activation energy. It should be noted that, in general, the converse would also be true: if the

conversion function was known, a single non-isothermal experiment would yield the kinetic parameters.

Unfortunately, additional complexity surrounds the use of non-isothermal methods. If the activation energy is not known with a high degree of accuracy, a non-isothermal data set may lead to an “apparent” conversion function [260]. This function will be distorted from the true form due to the systematic error introduced by the temperature ramp and the incorrect activation energy used to calculate the conversion function. This is caused partly by the mutual correlation between the activation energy and the pre-exponential factor [261] and is generally known as the kinetic compensation effect (KCE).

The KCE is caused by a variety of factors, which may include true physico-chemical changes in the sample, experimental procedure and mathematical artefacts [261]. In most cases this effect is characterised by a logarithmic relationship between the activation energy and the pre-exponential factor [261]. This form is attributed to the mathematical procedures used which allow an increase in one parameter while offsetting this by increasing the other [262], thus implying that this effect has no chemical significance.

Numerous approaches have, however, been proposed to compensate for it. Agrawal [263] suggested using a combination of multiple heating rates and isoconversion methods which are then compared with the integral method to yield self-consistent results. An alternative is suggested by Lyon [264] using a slightly different approximation of the temperature integral. Both these methods and several others claim to be capable of determining the activation energy independently of any assumptions regarding the functional form of the conversion function.

However, these methods were shown to be erroneous since they are all variations of the Coats-Redfern approach. This expression depends on the effective exponent, i.e. an implied conversion function, present in the equation which, in turn, influences the activation energy determined [262]. An alternative approach was recently suggested by Galwey [265] using at least two non-isothermal experiments at different scanning rates. The reaction rate is then calculated as a function of temperature for a range of constant increments of the conversion. By using these data, the activation energy for each of these intervals can be calculated. Thus the behaviour of activation energy as a function of

conversion is established and can be assessed. Finally, a pseudo-isothermal curve can be constructed and analysed in the manner previously discussed.

A related issue that has been the subject of much contention in the recent literature is the concept of a “variable activation energy”. This concept was probably first popularised by Vyazovkin [266];[267]. The motivation for a variable activation energy is based a variety of effects such as: the fact that for some reactions the free energy barrier of a chemical step is at least partially determined by the physical properties of the medium; in some cases diffusional limitations develop as the reaction proceeds; and, in general, kinetic processes are multistep processes.

This has led to a renewed focus on understanding the concept of activation energy, as originally proposed for homogeneous chemical kinetics, in terms of heterogeneous chemical kinetics [235];[268] as mentioned earlier. However, due to the empirical nature of this proposition, it is considered a retrograde step [269] which does not extend the theoretical scientific foundations on which current analysis is based. Alternative approaches to understanding the activation energy of solid state reactions have been proposed.

According to the L'vov theoretical model [270], the primary step in a solid state decomposition is the congruent dissociative evaporation of the reactant to give the primary products. Through equilibrium considerations, this “physical” approach provides an alternative to the “chemical” Arrhenius approach. However, there are a variety of reasons why this approach has not gained widespread acceptance [268]. These are, among others, the lack of incorporation of the interface geometry and the fact that the analysis is based on a single measurement; the approach has as yet had only limited application to very specific types of reaction.

It is interesting to note that a similar concept has been proposed by *Hurt et al.* [110] to account for the loss of reactivity for coal combustion at high conversions. It is speculated that the loss of reactivity is caused by a variety of factors, including thermal annealing of active sites, ash inhibition or preferential consumption of more reactive species. The authors note that the loss in intrinsic reactivity may be attributed to either the activation energy or the pre-exponential factor. However, due to a reported narrow range for the observed activation energies, this effect is incorporated into the pre-exponential factor. Despite this

fairly superficial treatment, especially regarding the deep interrelationship between these two parameters, this study does allude to the development of a deeper understanding of carbon oxidation based on proper mechanistic understanding.

Due to the complexity of the problem under consideration, significant questions have arisen around repeatability. Experimental results for identical samples from different laboratories, conducted on different thermal analysis machines, may deliver widely varying results. To this end several projects [271];[272] were launched to compare and standardise the approach to obtaining consistent kinetic data. Issues such as the physical nature of the sample, representative sampling, kinetic method, secondary effects and experimental set-up were considered. The results indicated that, despite the complexities, it is possible to get representative, consistent results if the experimental methodology is sound.

From the numerous publications mentioned here, it becomes clear that to gain any useful significance from experimentally determined kinetic parameters, especially activation energy, detailed physical and chemical characterisation of the solid reactant is required. This includes a thorough understanding of the microstructures, textures, crystal structures, imperfections, strains and participating components in the reaction. Preferably, complementary techniques should be used to determine physical parameters that can be directly linked to the mechanism being considered.

Furthermore, linking the shape of the reaction rate curve should not be considered a method for determining the mechanism. Instead, a clear evaluation of the physical parameters and structural understanding gained from the aforementioned considerations should be used to formulate a mechanism based on real, physically observable phenomena. In this fashion a tangible link between the mechanism and the material under consideration is ensured. This will allow expansion of the understanding to include other similar materials, as well as the incorporation of secondary effects caused by the presence of impurities.

A factor that is rarely taken into account in kinetic evaluations is the effect of the particle size distribution (PSD). The effects of the PSD on the three-dimensional phase boundary and Ginstling-Brounshtein diffusion-controlled reaction mechanisms are illustrated by Koga and Criado [273]. The investigation

is extended to include other variations of these models and a few pertinent conclusions are reached [274]. Firstly, the activation energy can be determined from isothermal analysis without the need to consider the PSD. However, the isoconversion methods can provide an activation energy that is independent of the PSD for all types of data. Furthermore, by using the kinetic curves at infinite temperature, the effect of the PSD can be incorporated to allow the correct evaluation of the pre-exponential factor and the kinetic model.

Very few studies have focused on using the solid state kinetics approach to modelling the oxidation of carbon. This is probably due to the myriad microstructures that are possible and the extensive surface heterogeneity. Burnham and Braun [275] have applied this approach to a variety of materials, including synthetic polymers, petroleum and oil shales, as well as coal. However, they find it necessary to utilise the logarithmic relationship between the pre-exponential factor and the activation energy distribution to account for the observed behaviours. This expression is identical to the one previously noted for the KCE, which calls the validity of this approach into question.

Pakula and co-workers [276];[277] suggested the use of the nucleation and growth mechanism to account for basal plane pitting during the oxidation of HOPG. Regrettably, the mechanism is never analytically linked with experimentally measured data or kinetic parameters. On the whole, no set of conversion functions and corresponding mechanisms have been identified for graphite. A recently published study by the author of this thesis, and others, on natural graphite [278] indicates that in order to understand the observed conversion functions of graphite, a more detailed mechanistic investigation is needed. However, good correspondence between isothermal and non-isothermal techniques was observed. This gives a preliminary indication that perhaps surface heterogeneity and other temperature-dependent effects may not pose insurmountable challenges for graphite as solid reactant.

The previous discussions also create concerns surrounding the work of Li and Brown [222] mentioned earlier. They assume a power law model for the conversion function of coal – one of the models with a very weak theoretical basis – with no analysis of the microstructure or even the particle size distribution. The lessons learnt from the solid state kinetics approach suggest that perhaps a

re-evaluation of the observed data for one of the most complex solid reactants, carbon, is called for.

2.13 Microstructure of graphite

As can be seen from the previous section, the reaction kinetics are intimately connected with an understanding of the geometric and physical structures of the solid, so it is prudent to evaluate the microstructure of graphite materials. From the preceding discussions it should be clear that the microstructures of natural and synthetic graphite materials would be vastly different. From an oxidative perspective it is perhaps more pertinent to ask: in which ways are they the same? In general, most oxidative studies on carbon materials in terms of microstructure focus on the ASA. This is a direct consequence of the relatively small crystallites present in these materials; in other words, the active sites are so prolific that they overwhelm crystallographic considerations.

However, for highly crystalline, natural graphite materials this is not the case. Hedley and Ashworth [279] studied the imperfections in Madagascan flake graphite. They found a high incidence of stacking faults and dislocation in the basal plane, directly observable by TEM. Unfortunately, much of this early work is intangible in modern digital form since the scanned versions of the original images are little more than dark grey-black smudges. More recent work by Atamny and Baiker [280] using STM clearly highlights the variety of microstructural defects possible in graphite.

As mentioned previously, recent studies on HOPG [195];[196];[197] have shown that for these highly crystalline materials oxidation proceeds mainly from lattice defects and results in the formation of pits in the basal surface. Inevitably, the samples were oxidising at their edges, but it was observed that the defects act as low-energy activation sites for oxidation, which is consistent with crystallographic considerations of active sites. In general, as can be seen from the work of Hongsheng *et al.* [30] and Bahattacharya *et al.* [226], both natural and synthetic graphite powders have particles with roughly flake- or disc-like shapes, i.e. flat particles with a small thickness and a relatively large diameter or width.

A perfectly flat basal plane based on optical observation can still have a high reactivity if the defect concentration is large and oxidation proceeds from these points. Another fact to bear in mind is that the presence of these defects would destroy the flake-like behaviour of a given sample since the sample will oxidise not only from the edges, but also from a random distribution of pits within the greater disc-like structure. This dramatically affects the active surface area development of such a shape. Furthermore, most naturally occurring flakes, while having a rough disc-like shape, also have structural detail in the form of folds, mosaic structures and agglomerates, as shown by the close examination of highly crystalline Ticonderoga natural graphite by Kavanagh and Schlögl [281]. This further complicates the development of active surface area in these materials.

As mentioned previously, from a kinetic perspective it is not enough to know the basic geometric shape and structural morphology of the particles. The PSD should also be taken into consideration. This fact is sometimes neglected in kinetic studies but is highlighted by the results of Jiang *et al.* [282] which clearly shows the relationship between particles of different sizes and their reactivity. Larger particles are expected to have a lower ASA per mass and hence a lower reactivity, leading to a decrease in ignition temperature. This work is substantiated by Zaghbi *et al.* [283] who show the relationship between BET surface area and particle size for natural graphite flakes.

Ultimately, this raises the issue of the extent of graphitisation of synthetic materials compared with natural materials. Within the porous structure of these materials the clear, flat layered surfaces characteristic of graphite are clearly noticeable, as can be seen in the work of Lim *et al.* [284]. However, it is unclear whether the crystallinity of these structures can be considered to be of the same degree as that of the visually equivalent natural graphite material. In some cases these structures are several microns across, which seems to be in contradiction to the crystallographic parameters for these materials measured with XRD and Raman since crystallite sizes for these materials are not expected to be more than 100 nm [24];[25].

In cases where the layered nature of the particle is not readily visible, understanding the microstructure of the synthetic graphite becomes even more complex. Recall the structures proposed for amorphous carbon materials

proposed by Franklin, Ban *et al.* and Crawford and Johnson in Figure 2-43, Figure 2-44 and Figure 2-45 respectively. Visually discerning between the crystalline structures and the amorphous interconnecting regions becomes problematic. It is impractical to use TEM techniques to examine the structure since preparing sufficiently thin samples without damaging the structure in the process is very difficult.

Attempts have been made to model this structure using a multiscale, porous, microstructural model [285];[286]. The model subdivides the large (micron-sized) granular structure into individual particles with large voids. The granules are then further subdivided into individual crystallites separated by microvoids. The work was successfully applied to the diffusion of atomic hydrogen in graphite for fusion applications [287] and shows that the degree of disorder of the internal surfaces and the closed porosity of the graphite are key considerations.

An even more relevant factor to consider is the expected development of this microstructure as oxidation proceeds. Acoustic microscopy has been used to measure the density change of oxidised nuclear graphite in order to explore this development [288]. However, this approach can only show a density profile along a single axis of a bulk graphite sample and the measurement carries a high degree of uncertainty.

A recently developed technique that shows a lot of potential in this regard is X-ray tomography [58];[59]. This technique irradiates a sample with X-rays while rotating the sample through 180° in equally spaced steps. In this way the local density and porosity of the material are measured non-destructively in three dimensions. Using the finite element-based approach, this technique has been applied to model the effect of oxidation on structural failure through porosity development [289]. Unfortunately, the technique has not yet been applied to model porosity development as a function of oxidation and, as mentioned previously, it still has limited spatial resolution.

An interesting approach to modelling the pore development in char samples has been suggested by Juniprom *et al.* [290]. They assume that the char consists of bundles of parallel graphitic layers. Using a random number generator, the oxidation of these layers is allowed to proceed randomly. It is stated that this is a representation of the degree of heterogeneity of each layer.

This leads to the development of micropores during the gasification as some layers are preferentially removed. These researchers found reasonable agreement between the model and the development of surface area and pore volume for a Logan seed-derived char.

A further complication arises during the processing of the samples to obtain powders for kinetic studies. Graphite is fairly malleable and can easily be transformed into amorphous carbon via milling [69]. Unfortunately, in the case of both natural and synthetic materials, it is unavoidable that solid samples must be milled at some point during the production of a powdered sample. It is inevitable that during this process complex structural changes are induced in the graphite. A remarkable point to note is that milling under different atmospheric conditions results in different microstructures. Milling in oxygen was found to suppress the fracture rate, generating crystalline samples with low surface areas, but largely anisotropic with many stacking faults [291]. A similar effect was found for milling under hydrogen atmosphere [292].

It should be noted though that all the approaches mentioned are aimed at modelling the bulk behaviour or the basal plane development of individual pits. None of the studies has focused on a single particle in its entirety. Understanding the microstructural effects, especially those of synthetic graphite materials, is clearly complex. Given the number of studies on the oxidation of graphite that have focused on characterising the kinetic parameters, it seems absurd that none of them has focused on first firmly understanding the development of the microstructure during oxidation.

2.14 Catalytic oxidation

A subject that has given rise to almost as wide a proliferation of literature as the oxidation of carbon is the catalytic oxidation of carbon. This reaction is of particular industrial interest as increasingly more strenuous demands are placed on the efficiency and throughput maximisation of our existing energy infrastructure. The fuel most widely used for this industry is coal and its related materials, which inevitably occur naturally with a wide variety of inorganic impurities. The catalytic rate-enhancing effects of these substances have been known for a long time [293]. Unfortunately though, despite extensive studies, the

exact mechanism surrounding these effects is still largely speculative. This section tries to address some of the myriad observed behaviours and compositional combinations possible for the catalysed reaction.

2.14.1 Metallic catalyst categories

As a starting point, it is convenient to classify the types of metallic catalyst according to their position on the periodic table. A justification for this would be the fact that the electronic structure of a catalyst plays a role in the catalytic mechanism, as will be seen in subsequent discussions. At this early point it is critical to take note of the fact that the catalytic effect of additives on the carbon oxidation rate is noticeable at trace levels, in some cases below 1 ppm [293];[294]. This makes independent study of the effects caused by different substances very difficult on samples that are not of the utmost purity.

Alkali metal salts are among the oldest known additives that markedly increase the reaction rate of carbon oxidation and, in general, carbonates are more active than other salts such as sulphates or halides [295]. In early studies it was noted, however, that a saturation effect was achieved when adding potassium carbonate to coke. Above 2%, further addition of the salt did not decrease ignition temperature by much. The relative reactivities of the alkali carbonates were reported by McKee and Chatterji [296] as Rb, Cs, Li > K > Na. Furthermore, this study showed that the carbonates of these metals generally all formed highly erratic, random channels on graphite. From *in situ* electron microscopy, the particles seemed to resemble liquid droplets, despite the reaction temperatures being at least 200 °C below the melting point of the carbonate. To explain this effect the authors postulate the existence of an oxidation-reduction cycle that occurs at the salt-metal interface, leading to the formation of a peroxide or other higher oxide, which has a significantly lower melting point.

2.14.1.1 Barium

Among the alkaline earth metals, barium exhibits a comparatively strong catalytic effect, while calcium is relatively weak. Although there is some discrepancy, the reactivity may be ranked, according to McKee [295], as: Na (comparatively) > Ba > Sr > Mg > Ca. However, the author remarks that

differences in the dispersion of the catalyst on the substrate may be a causal factor for the variation in reactivity measured in different studies. Barium, deposited as barium nitride on graphite, was observed [297] to exhibit erratic, random channelling behaviour below 750 °C. However, above this temperature particles were observed to shrink as material was deposited on the channel walls due to a wetting and spreading action. These thin films continued to gasify the graphite, leading to a saw-toothed or fluted appearance, with edge recession parallel to the armchair or $\langle 10\bar{1}0 \rangle$ crystal planes. It is interesting to note that in the same study, the barium did not exhibit wetting and spreading behaviour in a hydrogen environment. This is attributed to the fact that an oxide is not readily reformed following oxidation-reduction at the interface.

2.14.1.2 Vanadium

A large collection of studies have focused on the catalytic effects of the transition metals. Generally, vanadium is one of the more reactive species. Baker *et al.* showed [298] that the effects of metallic vanadium and vanadium pentoxide were similar, suggesting the same active species in both systems. The catalyst particles were liquid-like and propagated channels in the graphite with hexagonally faceted tips. The particle sizes and channel widths became progressively smaller as the channels increased in length, indicating a wetting effect similar to that noted for barium previously. It was observed that some material appeared to penetrate between the graphite layers. However, the exact nature of these formations is unclear as they were found to move freely without causing strain in the graphite. This indicates that they were on the surface of the graphite and it may be that, due to the thin samples used for this TEM work, the particles were actually moving along the bottom of the sample rather than being intercalated.

2.14.1.3 Molybdenum

A similar effect was reported by Baker *et al.* [299] for molybdenum. However, in this case differences were observed for pure molybdenum compared with molybdenum trioxide. For molybdenum trioxide both pitting and channelling were observed. The pitting produced hexagonal holes as the catalyst penetrated

through the basal plane and the pits expanded solely by uncatalysed reaction. The pits were observed to be parallel to twinning bands, indicating that they were bound by armchair faces. However, in some cases as the pits deepened, they were found to twist through a 90° angle and become progressively more circular in shape.

Channel propagation was found to occur if particles were located at steps or edges of the basal plane. Particles in the pristine regions of the basal plane remained catalytically inactive but were found to be mobile. Furthermore, in some cases channels were observed to become perceptibly deeper during propagation and, although the channels showed no particular orientation, the larger catalyst particles were hexagonally faceted. It was noted that small particles cutting shallow channels gave the highest reaction rates. Metallic molybdenum showed similar pitting behaviour, but the channelling showed definite orientation with respect to the underlying graphite, i.e. it was parallel to the armchair edges. Again, smaller particles were found to exhibit faster oxidation rates.

2.14.1.4 Metal oxides and carbides

The oxidation behaviour of the oxides and carbides of tungsten, tantalum, chrome and molybdenum was studied by Yang and Wong [300]. Among all the compounds considered, only molybdenum trioxide exhibited channelling, in accordance with the previous results, while all the others showed pitting. For carbides, the pits were hexagonal in shape but orientated parallel to zig-zag edges. The edge-recession rates measured were significantly higher than those for the uncatalysed reaction, despite the fact that the edges of the pits were distant from the catalyst particle. The authors concluded that the carbide particles act as dissociation centres, with subsequent spill-over of the oxygen and migration to the edges for reaction. For oxide particles, the pits were roughly circular.

2.14.1.5 Nickel

Nickel metal has been shown to be very inactive as a catalyst in an oxygen atmosphere [301], exhibiting only minor pitting activity at high temperature and some channelling at very high temperature (>1 065 °C).

In hydrogen, however, channelling was found to commence at temperatures as low as 845 °C. Channels show predominantly orientation parallel with the zig-zag edges, with hexagonally faceted channel tips. It is interesting to note that in a subsequent, study Baker *et al.* showed [302] that when alloyed with iron, the nickel-iron composite demonstrated behaviour that was quite different from that of either of the pure components. Iron on its own is found to induce both pitting and channelling in the graphite surface [295].

2.14.1.6 Chromium and copper

An effect similar to that observed for the metal carbides mentioned previously was found by Baker and Chludzinski [303] for chromium. Catalyst particles remained motionless but induced a measured acceleration in the edge-recession rate all over the sample. The authors also deduce the presence of an oxygen spill-over mechanism. The same study found that copper initially formed channels, which diminished at higher temperatures due to particle spreading and edge wetting. This is presumed to occur due to the formation of copper oxide. McKee [304] also studied the effect of copper and concluded that copper oxide particles undergo an oxidation-reduction cycle, similar to that discussed previously. The channels were found to be randomly orientated and when alloyed with zinc, aluminium or tin, the catalytic effect was inhibited.

2.14.1.7 Silver

Harris and Feates [305] observed that silver metal is a very active catalyst for carbon gasification. Under slow heating conditions, the catalyst forms pits and channels not located at edges or steps. It is postulated that this is due to direct removal of carbon from the basal lattice rather than vacancies already present. This conclusion was reached by comparing the number of catalytically active particles with the natural vacancy concentration of the graphite crystals. At higher heating rates, channelling was initiated only at steps in the graphite. Despite the channel tips having faceted particles, the channelling direction was found to be irregular.

2.14.1.8 Platinum and palladium

The effect of platinum and palladium on the oxidation of graphite was also studied by Baker and co-workers [306]. They found very similar behaviour for these metals: at low temperatures pitting was favoured, whereas at higher temperatures channelling became dominant. Particles that were active tended to have a spherical shape, while particles that remained stationary were irregularly shaped. They concluded that at the interface of the particle and the carbon, where gasification is occurring, the temperature is higher than in the bulk material. The channel propagation rate was found to be inversely proportional to particle size and channel depth, indicating that all active particles, irrespective of size, gasify the same amount of carbon per unit of time.

2.14.1.9 Lead

Lead has been shown to be one of the most catalytically active metals studied [295]. At low temperatures lead shows very high edge-recession rates of the graphite basal planes, followed by catalytic channelling and cessation of rapid edge recession. Harris *et al.* postulate [307] that due to the low melting point and wettability of graphite by lead, a very thin film of lead initially occupies the entire edge. As the oxidation proceeds, the edge lengthens and recedes, so that the film spreads out until it breaks and small beadlets of catalyst proceed to cut channels into the graphite. At high temperatures rings of irregularly shaped particles appear, probably due to lead decoration of expanded vacancy pits.

2.14.1.10 Boron

The catalytic effect of boron remains something of a mystery. It has been shown by Roscoe [308] that boron probably forms boric acid in moist oxygen, which directly chemisorbs at lattice imperfections. This complex is thought to be the catalytically active species, although the mechanism of catalysis is unclear. Furthermore, it was found that in dry oxygen boron inhibited oxidation. It is thought that this is due to the formation of boron oxide, which tends to block the active sites. The duplicitous action of boron warrants further study to understand its role in carbon oxidation [295].

2.14.2 Related reactions

As mentioned in the previous section, some catalysts showed marked differences in their catalytic activity and behaviour in different gaseous environments. In an effort to unravel the catalytic mechanism at the reaction interface, it is prudent to evaluate some of the observed behaviours in different gases. Tomita and Tamai [309] studied the effect of various transition metal catalysts on the hydrogenation of a variety of carbons. They found that the relative reactivities were almost exactly the reverse of the sequence found for carbon oxidation. The authors stress that the reactivity is strongly dependent on the preparation method of the catalyst-carbon system, especially regarding dispersion and the nature of the catalyst-carbon interface.

In a subsequent study on the catalytic hydrogenation of graphite by a variety of transition metals [310], the same authors show that both irregularly shaped channels and straight channels along preferred crystallographic directions, zig-zag in this case, are possible. They observe that the probability of curving channels being formed is found to increase with increasing particle size. Furthermore, an irregularly shaped particle was found to move so as to maintain the largest contact area with the peripheral face of the graphite. From visual inspection of the micrographs, smaller particles appear liquid-like and tend to form straight channels, while larger particles are usually erratically shaped and form curved, random channels. Tomita and Tamai propose an atomistic model based on different reactions at the graphite edge sites and the relative reactivities of the zig-zag and armchair faces. Larger particles are postulated to have a wide variety of contact points with the graphite surface, each with varying reactivity, as opposed to smaller particles. Hence the latter are far more likely to maintain a constant channel direction. It is also critical to note that their results suggested that, in general, the larger the particle, the faster the channel-propagation rate. This is the opposite of what was stated for oxidation earlier, indicating different rate-controlling steps for the two reactions.

Keep *et al.* [311] conducted similar studies on the nickel-catalysed hydrogenation of graphite, but taking critical care to consider channel depth and particle contact area. They found that the carbon gasification rate was not proportional to channel width, but rather to particle surface area. This implies that

the rate-controlling step is not the reaction at the catalyst-carbon interface. They propose that the rate-controlling step is the chemisorption or dissociation process of the reactant gas on the catalyst particle. This is followed by bulk or surface diffusion to the catalyst-metal interface and rapid reaction. This is similar to the spill-over mechanism noted for transition metal carbides and chromium mentioned in the previous section. An alternative mechanism is proposed to explain the observed behaviour whereby carbon dissolution into the catalyst particle occurs, followed by diffusion to the surface and reaction. In this case the surface reaction is postulated as being rate controlling to maintain consistency with the particle surface area effect. Several observations in support of the previous mechanism are given, but it is shown that the latter is possible and future work should be done to confirm the correct mechanism.

Goethel and Yang [312] questioned the previous investigation on the basis that it is not possible to determine the channel depth accurately. Instead, they focused only on monolayer etch pits caused by the platinum-catalysed hydrogenation of graphite using gold decoration. They postulate that instead of the reaction anisotropy mechanism proposed by Tomita and Tomai [310], it is the differences in interfacial tension that are the key consideration, with the zigzag edge being the preferred wetting plane. Particle movement is thus claimed to be driven by adhesion forces between the particle and the graphite edge planes. Furthermore, the reaction rate is found to be zero order in hydrogen partial pressure, but still dependent on particle surface area. It is therefore concluded that the spill-over mechanism proposed above is invalid and the latter mechanism is the correct one, with surface reaction being rate controlling.

Using the gold decoration method mentioned earlier, monolayer etch pits generated on graphite by steam and carbon dioxide, using alkali and alkali earth metals, were studied by Chen and Yang [313]. They found that the channelling behaviour was essentially the same for both these gases. Furthermore, channelling was found to be independent of particle size for particles above a lower limit. Based on this, it was concluded that carbon-carbon bond breakage was the rate-limiting step for all metals studied by this group, except vanadium oxide. In an earlier study [314] it was postulated that vanadium oxide is not readily reduced to a more reactive, lower melting, higher oxide in gases other

than oxygen. Thus for this metal it is concluded that surface oxidation of the metal is the rate-limiting step.

For small particles it is postulated that the chemisorption of the gaseous reactants is decreased and hence the reaction rate is lowered. This work is extended [315] to include a variety of other gaseous environments, including hydrogen, oxygen and nitrous oxide. For all systems considered only carbon dioxide and nitrous oxide exhibit the formation of wedge-shaped channels. It is stated that under the reaction conditions, uncatalysed reaction rates are too low to account for this phenomenon and it is concluded that catalyst is deposited on the channel wall as the particle moves. As further support for carbon-carbon bond breakage as the rate-limiting step, it is noted that when an existing channel cuts through a monolayer step, the reaction rate does not change with the change in channel depth. Later, the same group of Yang and co-workers [316] also utilised molecular orbital calculations to substantiate carbon-carbon bond breakage as the rate-limiting step during reaction with nitrous oxide.

Konishi *et al.* [317] studied the effect of cobalt on the hydrogenation of HOPG. They found that small particles consistently produced channels with preferred orientation along the zig-zag edge. In some cases when these particles encountered a higher step, they would continue along their current trajectory and burrow underneath the top layer. Furthermore, the channels would undergo semi-periodic 60° changes in direction, thus continually maintaining the preferred movement along a zig-zag edge.

2.14.3 Topography of catalytic behaviours

As noted in the previous sections, pitting and channelling are commonly found during the catalysed oxidation of graphite. The topography of these behaviours and of catalyst mobility in general are still somewhat of a mystery and merit a deeper exploration. One of the earliest studies on these phenomena was done by Hennig [318]. He indicated that, in general, pitting behaviour is not commonly found on pristine basal planes. In fact, special treatments such as extensive mechanical deformation, quenching or particle irradiation were necessary to induce pitting. These treatments all induce large amounts of

vacancies into the crystal structure and it was concluded that pitting occurs only when catalyst attacks the graphite at vacancies or defects.

He found that all channels started at steps in the crystal surface and when located some distance away from a step, catalyst particles remained inactive. Straight and curved channels were observed for different types of catalyst. The straight channels showed orientation along the zig-zag direction in dry oxygen and along the armchair direction in moist oxygen, exactly the opposite to what was observed for uncatalysed oxidation. Hennig observed wedge-shaped channels and associated this with uncatalysed oxidation after the catalyst particle had passed. Hence a ratio of the half-width of the channel base to the channel length was used to calculate the ratio of uncatalysed to catalysed reaction rate. It was found that the ratio increased as a function of temperature and the activation energy for these two processes differed by 150 kJ/mol. Since the activation energy of the uncatalysed reaction was thought to be 190 kJ/mol, the catalysed reaction had an activation energy of 40 kJ/mol, but it is stated that there is large uncertainty in these calculations due to variations in channel dimensions. Catalyst particle agglomeration was also observed during extended oxidation.

Certain crystals were treated to induce pitting by heating to 2 700 °C and then rapidly quenching. During uncatalysed oxidation only a small number of random pits were noticed, but the number was slightly increased during oxidation in moist oxygen. However, when metallic catalysts were added, very extensive pitting was observed. This effect is confirmed by Harker *et al.* [319] who studied neutron-bombarded moderator graphite. It is interesting to note that Hennig observed that in damaged material any catalytic particle would form a channel, irrespective of whether it was in contact with a basal step or not. It was also observed that in this case the channels became progressively deeper, with the rate of penetration increasing with increasing defect concentration. When the channel depth approximated the particle radius, lateral motion ceased and pit formation commenced. In general, the formed pits were hexagonal in shape, except at high oxygen pressures when they became less regular.

A similar study was conducted by Presland and Hedley [320] using TEM. They, too, found preferred channel orientation along the zig-zag direction and wedge-shaped channels. In this case, though, catalyst particles not directly associated with steps were found to be mobile and capable of migrating rapidly,

especially under oxidising conditions, until an edge or step was encountered, at which point channelling commenced. These platinum particles behaved in a liquid-like manner around 900 °C, despite the melting point of platinum being 1 773 °C. There is evidence that this behaviour is dependent on the surface state of the graphite, as shown by heating in vacuum versus air, or when the graphite surface was coated with a carbonaceous film, which immobilised the particles.

As noted from the previous discussions, the details of particulate movement during catalytic oxidation are usually complex and vary from one catalyst to the next. Small particles are generally found to exhibit the most vigorous mobility, while particles above 10 µm usually become immobile and inactive [295]. Furthermore, in some cases the particles collide and coalesce, thereby losing mobility and activity, although in some cases this is also caused by gradual conversion of the catalyst to an inactive state.

In an effort to classify the wide variety of observed behaviours, McKee [321] analysed the catalytic effect of more than 27 metallic compounds on highly purified graphite. The metals were all introduced as acetates which underwent transformation to an oxide phase, the composition of which could be inferred from the thermo-balance measurement. He concluded that the observed topographical behaviours could be broadly categorised into five categories:

1. General erosion in all crystallographic directions (e.g. Pb)
2. Channelling:
 - (i) In preferred directions (e.g. V)
 - (ii) Irregular (e.g. Cu, Cd)
3. Etch pit forming:
 - (i) Hexagonal, zig-zag pits (e.g. Fe)
 - (ii) Irregular (e.g. Mn, Ag)

It is stated that in some cases, e.g. Mo, catalysts can attack by multiple modes. However, as noted earlier, some catalysts can even change mode as a function of temperature. Lead oxide was found to be the most vigorous catalyst and many oxide particles exhibited surface mobility, with irregular channelling being generally observed. McKee observes that pits induced by Fe have a zig-zag edge, whereas uncatalysed pits formed below 900 °C have an armchair edge. This finding is in accordance with the work of Thomas [10] which showed that zigzag edges have a slightly higher uncatalysed reactivity; hence the armchair configuration dominates. It is therefore also of interest to note that, in

general, catalyst channelling along a preferred direction occurs along the zig-zag edge, whereas one would expect armchair edges for uncatalysed oxidation. As stated by McKee, the reason for this inversion is still unclear, but based on the observations of other researchers, this effect appears to be dependent on the type of catalyst used. He also mentions that although the correlation is rough, metal oxides that are very stable and unlikely to be reduced do not tend to function as catalysts. This provides further evidence for the oxidation-reduction mechanism at the catalyst-particle interface.

A similar comprehensive study of metallic catalyst activity on graphite was done by Baker [322] using the “controlled atmosphere electron microscopy” (CEAM) technique. A specially designed reaction chamber is fitted into a TEM. This cell allows the reaction to be conducted at high pressure, while the TEM is maintained under vacuum. Furthermore, a platinum heating element and thermocouple allow precise control of the sample temperature up to 1 300 °C. In this way it is possible to monitor the catalysed oxidation of graphite *in situ*.

On the basis of several pieces of evidence, he concluded that particles with liquid-like behaviour or a faceted tip at low reaction temperatures were, in fact, molten due to much higher temperatures at the point where oxidation was occurring. However, this fails to explain the mobility of platinum particles on the basal plane, where no oxidation was occurring, at temperatures well below its melting point, mentioned earlier. Common features of all oxidation catalysts were found to be:

1. All channels originate from steps or edges.
2. The linear rate of channel propagation is the same in different crystallographic directions.
3. The width of the channel is generally the same as that of the particle generating the channel.
4. The linear rate of channel propagation is proportional to particle width and channel depth; hence small particles generating shallow channels have the highest propagation rate.
5. All channel propagation rates increase with increasing temperature.

Thus the conclusion mentioned earlier for platinum and palladium, i.e. that all particles gasify carbon at the same rate, is generalised to all metallic catalysts.

Furthermore, Baker [322] showed that there is excellent agreement between the bulk measurements made for the activation energy of a particular catalyst-carbon combination and the activation energy measured by CEAM. The uncatalysed activation energy is measured as 200 kJ/mol and a variety of catalyst activation energies are given, ranging from 70 to 400 kJ/mol. It is pertinent to note that this implies that the rate-controlling process for catalysed reaction controls the bulk reactivity.

In accordance with the assertion mentioned above, namely that due to high interfacial temperatures catalyst particles are molten, Baker in a later publication [323] tried to classify the observed catalytic activity on this basis. Three basic forms of attack are postulated, as shown in Figure 2-55.




Figure 2-55: Catalyst mode of attack (after Baker [323])

Firstly, for a non-wetting catalyst that has a contact angle of less than 90° , low reactivity will be observed and the catalyst will probably act via a pitting mechanism. Secondly, for a catalyst that exhibits a strong interaction with the graphite edge, atoms will undergo spreading to form a thin film. Such a catalyst will exhibit high activity and attack via the edge-recession mode. Thirdly, for catalysts that demonstrate an intermediate wetting state, the dominant mode of attack will be channelling by a discrete cap-shaped particle.

Based on the free energy data for oxide formation, the reactivity of several transition metal catalysts is successfully classified according to this scheme. It is stated that metals that readily undergo transition to an oxide will show a high degree of interaction with the graphite edge, inducing the spreading mode of attack. Metals that do not dissociatively adsorb oxygen tend to remain in their noble state and attack via the channelling mode. An attempt has very recently

been made by Datta [324] to model the mobility of liquid metallic particles on graphene. It is postulated that two possible factors explain the mechanism of etching: firstly, the difference between the equilibrium wettability of the graphene and the substrate that supports it; or, secondly and alternatively, the large surface energy associated with a graphene step – this mechanism is equally applicable where another graphene layer is the substrate, i.e. graphite in general. These factors lead to a difference in contact angle, which results in an unbalanced Young's force at the contact line that drives particle motion. While these approaches may be applicable for catalysts that undergo liquid-like transformations, some catalyst particles, i.e. those that retain their irregular shape and form random, erratically shaped channels, cannot be classified according to this framework.

Wu and Radovic [325] examined the effect of calcium and potassium acetates on the oxidation of a carbon/carbon composite. They found that catalytic effectiveness was dependent on pretreatment in an inert gas, catalyst loading method and amount, although the dependencies varied for the two types of catalyst. The catalyst reactivities were also found to be markedly higher for the fibres as opposed to the chemical vapour deposited matrix. It is postulated that the interfacial contact is fundamental to these observations. This indicates that the catalytic behaviour of the same catalyst will vary for different forms of carbon. Due to better wettability and higher mobility, potassium is a more effective catalyst and less dependent on the factors mentioned previously.

The effect of porosity on the catalytic oxidation process should also be briefly considered. Su and Perlmutter [119] applied the random pore model to coal samples impregnated with sodium and potassium carbonates. They found that catalyst penetration into the sample was fairly limited. Thus oxidation was proceeding by rapid, catalysed reaction at the particle exterior and uncatalysed reaction in the interior. Experimental results [326] indicated that despite the fact that the effect on the overall reaction rate was only a few fold, the intrinsic catalysed reaction rate was several orders of magnitude higher than the uncatalysed rate. The results further indicated that the catalyst served only to increase the number of active sites, but left the activation energy largely unchanged. This is probably due to the fact that a significant amount of the reaction was still proceeding via the uncatalysed route.

It seems very hard to rationalise all the observed topographical features of catalysed oxidation. The proximity, mode of contact and physical form of the catalyst particles and carbon substrate are of great importance in determining the magnitude and topographical details of the catalytic process [295], not to mention their reactivity. These effects are also intimately linked with the mechanism at the catalyst-carbon interface, which will be considered in more detail in the next section.

2.14.4 Catalytic mechanisms

The theories associated with the catalytic mechanisms mentioned in previous sections may be broadly classified into two categories [293]. Firstly, the oxygen-transfer mechanism regards the catalysts on the surface of the graphite as an oxygen carrier via the oxidation-reduction cycle. Thus the catalyst is assumed to undergo a continuous cycle between two oxidation states, e.g. metal to metal oxide, or lower metal oxide to higher metal oxide. This mechanism tacitly assumes that the rate-limiting step for the uncatalysed oxidation is oxygen chemisorption. The catalyst particle thus serves as an active chemisorption centre, consequently accelerating the reaction rate.

The second approach is the electron-transfer mechanism. This has a variety of interpretations, but involves mainly the transfer of electrons between the carbon and the catalyst, implying the establishment of a physical bond. This is because many of the known catalysts have unfilled energy bands that are capable of accepting electrons from the graphite matrix. Presumably, a redistribution of π -electrons in the graphite takes place, resulting in a weakening of the carbon-carbon bonds at the edge sites, while increasing the carbon-oxygen bond strength. As a justification for channelling along preferred crystallographic directions, Hennig [318] postulated the existence of these catalyst-carbon bonds.

A slight reformulation of these mechanisms, which is popular in the hydrogenation literature, is given by Tomita [327], namely the gas spill-over mechanism whereby the gas species adsorbs and dissociates on a metal catalyst. These atoms migrate through or on the particle to nearby carbon atoms and gasification occurs. The mechanism is shown schematically in Figure 2-56 and is the most popular mechanism for catalytic hydrogenation.




Figure 2-56: Spill-over catalytic mechanism (after Tomita [327])

The carbon dissolution mechanism is shown schematically in Figure 2-57 and is especially favoured in the case of nickel catalysts during hydrogenation. Carbon dissolves into the catalyst and migrates through or over the particle to the surface where reaction with the gas takes place.




Figure 2-57: Carbon dissolution catalytic mechanism (after Tomita [327])

These mechanisms are similar to the descriptions given by Walker and co-workers previously, with the emphasis on slightly different aspects. However, the rate-controlling process for these two mechanisms is ambiguous and can be carbon-carbon bond breakage (electron-transfer mechanism), or reactant gas chemisorption or dissociation (oxygen-transfer mechanism) for either of the two cases. Nevertheless, the inferred rate-controlling step is usually carbon-carbon bond breaking for dissolution and oxygen dissociation for spill-over.

Goethel and Yang [328] point out that an additional process could limit the reaction rate, namely transport of the atoms, oxygen or carbon, through the catalyst particle. The work of Baker and co-workers [306], mentioned earlier, stated that the linear rate of channel propagation was inversely proportional to particle size and channel depth. Based on this, Goethel and Yang concluded that in that case the diffusion through the catalyst particle was in fact the rate-limiting step.

If the channelling rate increased with particle size (at the same channel depth), i.e. an increase in the surface area for dissociation, it would indicate that the oxygen-transfer mechanism is rate controlling. If the channelling rate remained constant with an increase in channel depth, this would indicate that the carbon-carbon bond breakage was rate limiting. The case where an increase in the channel depth leads to a reduction in the channelling rate indicates that diffusion is rate limiting since the maximum resistance arises from the longer diffusion path to the bottom of the interface. If the particle is assumed to be spherical in nature, it may be similarly argued that when the particle size is increased, for the same ledge height, the diffusion path is also extended, leading to a reduction in rate.

Using only monolayer etch data, which avoids the uncertainty surrounding channel depth, Goethel and Yang [328] examined the rate-limiting step for the catalysed oxidation of graphite by Pt, V and Cu. For platinum the rate was found to be independent of particle size – much lower than a modelled diffusion-limited rate (discussed in Section 2.15) and higher than the rate observed for multilayer etching. Thus the rate-controlling step was concluded to be carbon-carbon bond breakage. For copper the channelling rate was again independent of particle size, but in this case it is unclear whether carbon-carbon bond breaking or the oxidation-reduction cycle of the copper at the interface is rate limiting. For vanadium the monolayer etch rate is proportional to the particle size and it is concluded that the gaseous oxidation of vanadium to vanadium oxide or of vanadium oxide to a higher oxide is rate limiting. This demonstrates subtle variations on the mechanisms proposed earlier, which indicates that they lack a holistic view of the phenomena taking place.

Of late, a growing interest in the use of small catalyst particles to cut graphene, due to its unique electronic properties, into controlled shapes for use in

field-effect transistors or as interconnects has emerged [329];[330]. An alternative approach has been suggested by Severin *et al.* [331] to use catalytic materials as the tip of a “catalytic pen” to do lithography on graphene in an oxidative environment. These applications have resulted in a resurgence of studies dedicated to elucidating the mechanism behind the catalysed oxidation of graphite. These studies have been no more successful than those done in the past in finding a single unifying mechanism, as pointed out by a very recent article by Biró and Lambin [332].

Schäffel *et al.* [333] found evidence for the spill-over mechanism for cobalt-catalysed hydrogenation, while Ci *et al.* [334] found evidence for the carbon-dissolution mechanism in nickel-catalysed hydrogenation. It is very interesting to note, though, that the latter study provides a plausible explanation for the consistent 60 or 120° turns observed during catalytic channelling along preferred crystallographic directions. It is postulated this occurs to maintain adhesive energy between the catalyst particle and the graphite edge. When a channelling particle encounters a step deeper than the channel itself, it will preferentially turn back at an angle of 60 or 120° to maintain channel propagation and adhesion. This has been proved experimentally [330];[332] and is thought also to occur when a particle encounters a lattice defect, such as a vacancy, as would be anticipated in a graphite crystal.

This and previous sections clearly illustrate the wide variety of catalytic behaviours that are possible, even for a single element. These are dependent not only on the oxidation state of the metal, but also on its physical state, i.e. liquid or solid. The situation is made even more complicated by the observation that particles of known composition exhibit liquid-like behaviour well below their melting points. This is a direct indication that interaction effects are present at the carbon-particle interface. Furthermore, some metallic catalysts show clear transitions between different oxidation states as a function of temperature. It is also apparent that the mechanism that is active at the carbon-particle interface, which ultimately determines the behaviour, is highly dependent on all of these factors.

Given the observed evidence in support of nearly all of these mechanisms, it seems highly likely that in fact all of these factors contribute in varying amounts to any given catalyst-carbon combination [293]. Thus it is a seemingly impossible

task to classify the catalytic mechanism and subsequent action of all possible impurities, even more so if the initial composition of the impurity is not known. In this case it may be prudent to pursue a more generalised approach based on the qualitative behaviour of the catalyst.

2.14.5 Modelling catalytic oxidation

An early attempt to model the kinetics of catalytic oxidation was made by Heuchamps and Duval [335]. They favoured the electron-transfer mechanism and based their model on the distribution of charges at the prismatic planes of graphite [293]. They proposed that the two reactions effectively occur in parallel: the uncatalysed reaction with a much higher activation energy (251 kJ/mol) and the catalysed reaction with a lower activation energy (153 kJ/mol). However, they found the marked presence of a catalytic compensation effect, exactly similar to the KCE discussed earlier for solid state reactions. In this case the relationship is as given below [295]:

$$\log(A) = \log(k_0) + \frac{E_A}{RT_S} \quad (2.58)$$

where T_S is the isokinetic temperature at which point all the reactions in the series proceed at the same rate.

Thus it is unclear whether the effect of the catalyst is to increase the pre-exponential factor, as would be expected through the creation of active sites, or to reduce the activation energy by alleviating some rate-limiting step, such as carbon-carbon bond breakage. An alternative approach was developed by Feates *et al.* [336]. They maintained the concept that two distinct reactions were occurring in parallel, but calculated the total rate constant as follows:

$$k = sA_C \exp\left(\frac{-E_C}{RT}\right) + (1-s)A_U \exp\left(\frac{-E_U}{RT}\right) \quad (2.59)$$

where A_C , E_C and A_U , E_U are the kinetic parameters of the catalysed and uncatalysed reactions respectively and s is the fraction of the carbon active surface area where the catalysed reaction is taking place

This presents the intriguing possibility of a compositional dependence for the overall reaction rate constant and hence the kinetic parameters, since the fraction of surface area under catalysed oxidation will inevitably increase as the particle shrinks. Furthermore, any increase in the initial catalyst concentration will tend to decrease the activation energy. However, since the kinetics of catalysed oxidation do not always exhibit this compensation effect, the validity of this model is not clear [295].

Choi *et al.* [337] used the carbon dissolution mechanism to derive a physical model for the concentration gradient of carbon within a single catalyst particle. Carbon is assumed to diffuse from the catalyst-carbon interface to the top surface where a thin viscous oxide layer forms. This layer is assumed to consist mostly of adsorbed oxygen and diffused carbon atoms; its thickness is constant. The model is capable of predicting the inverse proportionality observed for channelling rate versus particle size. Furthermore, it shows excellent correlation to the data of Baker and co-workers [306] for palladium.

A similar model was developed by Goethel *et al.* [338] based on their work mentioned earlier. It is also shown to be capable of accurately replicating Baker *et al.*'s data on channelling rate versus particle size, as well as experimentally observed rate constants. It is noteworthy that this model predicts, for a fixed channel depth, that the channelling rate variation with increasing particle size can be either decreasing or increasing, depending on the reaction rate constant. Thus the model is shown to be capable of predicting the data of Keep *et al.* [311] for catalysed hydrogenation, which is known to occur much more slowly than oxidation. These data showed an increase of channelling rate with an increase in particle size.

A rather simplistic thermal analysis route is chosen by Li *et al.* [339] to study the effect of three coal-burning additives on the oxidation of graphite. Using the empirical "order of reaction" model given in Table 2-12 and a single non-isothermal experimental data set for each combination, the kinetic parameters of the model are calculated. The reaction order (n) is found to vary nonsensically from 7.93 for pure graphite to values of 1.16, 1.41 and -2.46 for the

three additives. Furthermore, despite the fact that all three additives clearly show massive enhancements in the reaction rate as a function of temperature, the activation energies are determined to increase, remain roughly the same and dramatically reduce relative to that of pure graphite. This study clearly illustrates the dangers of blindly applying an empirical model for the reaction mechanism to measured thermal analysis data.

A completely unique approach to modelling the pitting and channelling action of catalysts in the gas-carbon reaction was presented by Ranish and Walker [340]. A graphite flake is represented by a single circular disc with discrete catalyst particles distributed evenly along the circumference of the disc. This is for the case of channelling catalytic particles, whereas for pitting catalysts the particles are uniformly distributed across the basal plane of the disc. The particles are assumed to burrow into the graphite in a direction normal to the exposed edge. Thus in the case of channelling particles, their movement is directed solely at the centre point of the disc. As the particles penetrate the disc, the pore walls are oxidised via the uncatalysed reaction. This is shown schematically in Figure 2-58.




Figure 2-58: Catalyst-induced pore geometry (after Walker [340])

To link the model with experimentally observed channelling behaviours, the ratio of the channel length divided by the channel exterior half-width, κ , is used. As mentioned in Section 2.14.3, this parameter represents the ratio of the catalysed to uncatalysed reaction rate and was found to vary from 15 to 80. As the catalyst particles penetrate the disc and the pore walls grow, a point will be

reached where the pore mouths on the circumference converge. At this point steady state roughness is achieved and active surface area will no longer increase as a function of burn-off.

Because the catalyst particles are randomly distributed along the circumference at varying heights, an assumption is needed to determine the point of steady state roughness. Ranish and Walker [340] assume that this occurs when the pore mouth width of the average number of particles in a circumferential strip as wide as the diameter of a catalyst particle is equal to the circumference of the particle. All values are thus reported as fractions of the steady state roughness achieved and, unfortunately, no attempt is made to model the active surface area development of this structure as a function of conversion.

For pitting catalysts the approach is similar with particles tunnelling inward, while the basal plane remains intact and the circumference shrinks. Assumptions are made to calculate the point of steady state roughness. Firstly, when steady state roughness is achieved, the areas of the pore mouths equal the original basal area and all the catalyst particles are within a pore radius from the edge. Secondly, in order to ensure that the catalysts are not being undercut by the receding flake circumference, the uncatalysed recession depth is set such that the reduction in basal plane area is less than 10%.

This model was used mainly to show that even at parts per million, the global reaction rate is heavily influenced by the catalysed reaction rate. This is not due to direct interfacial catalyst attack, but rather occurs by virtue of the increase in the active surface area caused by the catalyst. From an examination of the previous sections, a key factor can be highlighted: almost none of the investigations considered has focused on establishing the link between the creation of active surface area by the catalysts and the reaction rate.

2.15 Oxidation inhibition

A comprehensive review on the oxidation protection of graphitic carbon materials was performed by McKee [341]. In general, the inhibition of oxidation is thought to proceed via two possible routes: firstly, through the direct poisoning of the graphite active sites, and secondly, via the formation of an oxygen diffusion barrier. The protection of carbon/carbon composites presents a slightly more

complicated problem because of the differing materials, but the same principles apply.

Extremely extensive arrays of coatings have been tested for their oxidation inhibition properties. These include noble metal coatings, which carry high costs, and metal carbides, which lack clear performance data. Nitride coatings of aluminium nitride or boron nitride have shown some promise, but film orientation and boron oxide film formation are of great importance. Key considerations for an effective diffusion barrier are low oxygen permeability and a close match of thermal expansion to avoid crack formation.

In this regard silica coatings or multilayer materials, such as silica scales formed on silicon nitride or silicon carbide, probably offer the best approach for oxidation protection above 1 000 °C. This is because of the self-healing properties of silica above 1 500 °C and the fact that fused silica has the lowest oxygen permeability of any oxide. An additional surface glaze, possibly a boro-silicate, may help to heal incipient cracks during thermal cycling. However, the pressure build-up of gaseous reaction products may pose a problem for long-term applications. Other multilayer coating systems that show oxidation inhibition potential require further development to match the thermal expansion coefficients of the different layers.

Boron, in the form of borate glazes, has been shown to offer oxidation protection for carbon materials. Carbon powders sintered with boron carbide and silicon carbide to form a composite also show a large reduction in the oxidation rate. Direct boron substitution into the graphite matrix may also lead to the formation of strong carbon-boron bonds, which increase oxidation inhibition [318];[342]. However, as mentioned earlier, the oxidative behaviour of boron is complex and requires a deeper understanding to be fully utilised.

The inhibitive effect of phosphorous has also been used for a long time to protect carbon materials. Phosphate glazes can be used, often mixed with borates, as a protective coating. However, treatment with phosphoric acid alone (or in combination with boric acid) has been reported to lower the reactivity of carbon with air at elevated temperatures. Halogens have been found to be successful in reducing the oxidation rates of carbon when added in small concentrations to the oxidising gas. In combination with phosphorous

compounds, e.g. phosphorous oxychloride, much stronger poisoning effects are observed.

A recent study by Wieber *et al.* [343] concluded that the inhibition of oxidation by phosphoric acid was caused by active site blockage, leading to only a small reduction in the activation energy. Furthermore, the effect was found to be significantly reduced above 800 °C, presumably due to evaporation of the phosphate complex. It is interesting to note that addition of phosphoric acid to samples contaminated with alkali metal salts resulted in a reduction in reactivity. This is presumed to occur due to neutralisation of the salts by the acid.